

AD _____

Award Number: W81XWH-05-1-0041

TITLE: Prostate Dose Escalation by Innovative Inverse Planning-Driven IMRT

PRINCIPAL INVESTIGATOR: Lei Xing, Ph.D.

CONTRACTING ORGANIZATION: Stanford University
Stanford, CA 94305-5401

REPORT DATE: November 2006

TYPE OF REPORT: Annual

PREPARED FOR: U.S. Army Medical Research and Materiel Command
Fort Detrick, Maryland 21702-5012

DISTRIBUTION STATEMENT: Approved for Public Release;
Distribution Unlimited

The views, opinions and/or findings contained in this report are those of the author(s) and should not be construed as an official Department of the Army position, policy or decision unless so designated by other documentation.

REPORT DOCUMENTATION PAGE

*Form Approved
OMB No. 0704-0188*

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing this collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. **PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.**

1. REPORT DATE (DD-MM-YYYY) 01/11/06	2. REPORT TYPE Annual	3. DATES COVERED (From - To) 1 Nov 2005 – 31 Oct 2006		
4. TITLE AND SUBTITLE Prostate Dose Escalation by Innovative Inverse Planning-Driven IMRT		5a. CONTRACT NUMBER		
		5b. GRANT NUMBER W81XWH-05-1-0041		
		5c. PROGRAM ELEMENT NUMBER		
6. AUTHOR(S) Lei Xing, Ph.D. E-Mail: lei@reyes.stanford.edu		5d. PROJECT NUMBER		
		5e. TASK NUMBER		
		5f. WORK UNIT NUMBER		
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Stanford University Stanford, CA 94305-5401		8. PERFORMING ORGANIZATION REPORT NUMBER		
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) U.S. Army Medical Research and Materiel Command Fort Detrick, Maryland 21702-5012		10. SPONSOR/MONITOR'S ACRONYM(S)		
		11. SPONSOR/MONITOR'S REPORT NUMBER(S)		
12. DISTRIBUTION / AVAILABILITY STATEMENT Approved for Public Release; Distribution Unlimited				
13. SUPPLEMENTARY NOTES				
14. ABSTRACT: The goal of this project is to develop innovative inverse treatment planning techniques for prostate radiation therapy. In the last funding period significant progress has been made toward the goal of the project. We have established the voxel-based dose optimization framework for IMRT treatment planning. Realizing that current radiation therapy does not adapt to inter-fraction organ movement and dosimetric errors caused by inaccurate patient setup or organ deformation during a course of treatment, we have devised a dynamic closed-loop control algorithms for adaptive therapy (ART) inverse planning and demonstrate their utility with data from phantom and clinical prostate cases. The inverse planning for ART is particularly relevant considering that the emergence of on-board cone beam CT in the clinics, which affords an effective means to obtain the patient's geometric model just before treatment and recompute on a routine basis the dose to be delivered (or actually delivered) to the patient. Our work makes it possible to adaptively taking the organ deformations and dose delivery history into account. We anticipate that these tools will greatly facilitate the imaging, planning, delivery, and quality assurance of prostate IMRT.				
15. SUBJECT TERMS Prostate Cancer				
16. SECURITY CLASSIFICATION OF: a. REPORT U		17. LIMITATION OF ABSTRACT UU	18. NUMBER OF PAGES 209	19a. NAME OF RESPONSIBLE PERSON USAMRMC
b. ABSTRACT U				19b. TELEPHONE NUMBER (include area code)
c. THIS PAGE U				

Table of Contents

Introduction.....	4
Body.....	4
Key Research Accomplishments	7
Reportable Outcomes.....	7
Conclusions.....	9
References.....	9
Appendices.....	11

I. INTRODUCTION

This Idea Award (PC040282, entitled "Prostate Dose Escalation by Innovative Inverse Planning-Driven IMRT") was awarded to the principal investigator (PI) for the period of Nov 1, 2004—Oct. 31, 2007. This is the annual report for the second funding period (Nov. 1, 2005 – Oct. 31, 2006). The goal of this project is to improve current prostate IMRT by establishing a novel inverse planning framework. Under the generous support from the U.S. Army Medical Research and Materiel Command (AMRMC), the PI's research team has made significant progress toward the general goal of the project and contributed greatly to prostate cancer research. A number of significant conference abstracts and refereed papers have been resulted from the support. The preliminary data obtained under the support of the grant has also enabled the PI to start new research initiatives. In this report, the past year's research activities of the PI are highlighted.

II. RESEARCH AND ACCOMPLISHMENTS

Inverse treatment planning is at the foundation of modern radiation therapy. We have worked in the past year to improve the existing inverse planning framework. Our work has been focused on four related areas in inverse planning, which are: (1) inverse planning strategies for adaptive prostate radiation therapy; (2) relationship between voxel-based penalty scheme and EUD framework; (3) enabling techniques for clinical implementation of adaptive prostate radiation therapy; and (4) evaluation of adaptive inverse planning for prostate IMRT. By using the new inverse planning techniques, we can now significantly improve the radiation dose distributions as compared with the current practice. The study should have widespread impact on clinical prostate IMRT in the future.

Inverse planning for adaptive prostate radiation therapy: Current IMRT treatment plan optimization and dose delivery are two decoupled steps. In each fraction, the patient geometry is hardly the same as in the pre-treatment CT simulation. A commonly used method to take the uncertainty into account is to add a safety margin, whose size is based on population statistics, to the target and sensitive structures. This significantly compromises the success of radiation therapy. Recently, on-board cone-beam computed tomography (CBCT) integrated with a medical linear accelerator has become available and promises to improve the situation. CBCT provides a valuable 3D (or even possibly 4D¹⁻³) geometric model of the patient in the treatment position. This not only affords an opportunity for on-line correction of patient setup error^{4, 5}, but also makes adaptive radiation therapy (ART)⁶⁻⁸ possible, which uses the volumetric information to adjust the treatment plan each fraction according to the updated patient anatomy and

positioning. ART can optimally compensate various uncertainties, including organ deformation and inter-fraction organ motion as well as dosimetric errors incurred in previous fractions⁹⁻¹³. To realize ART clinically and maximally exploit the potential of this new form of image guided radiation therapy (IGRT), a robust inverse planning strategy for ART must be in place.

We have recently developed dynamic closed-loop control strategies for ART inverse planning and demonstrate their utility with data from phantom and clinical cases¹⁴. Closed-loop control algorithms are a general tool for dealing with time-dependent systems^{15, 16}. The algorithms in all these applications share the same basic closed-loop control framework of repeated re-evaluation and re-planning. ART is a natural application for closed-loop control because CBCT provides frequently updated system information. To meet different clinical requirements, we investigated two types of closed-loop control algorithms (*Adapting to Changing Geometry* (ACG) and *Adapting to Geometry and Delivered Dose*(AGDD)). ACG is useful when the accumulated dose is not known accurately. At this point, the deformable registration, which is essential for the calculation of accumulative dose, is still not robust enough for clinical use. ACG is particularly helpful in this situation as it affords a currently implementable technique to cope with the nuisance caused by organ deformation. The AGDD algorithm is designed to optimize the radiation treatment when both geometric and dosimetric updates are available from time to time.

We emphasize that the proposed ART inverse planning is purely dose based and doesn't consider any radiobiological effects. In principle treatment plan optimization should be based on biological models as they are clinically the most relevant. We have recently proposed a general time-dose-fractionation optimization strategy¹⁷. This idea combined with our ART framework could potentially handle biologically adaptive radiation therapy. This research is currently in progress.

Relationship of EDU-based and dose-based plan optimization: The essence of inverse planning is how to rank objectively the competitive treatment plans. We have established a unified inverse planning framework and shown the equivalence of EUD-based and the conventional dose-based objective function¹⁸. In the existing approaches, the dose-based function (e.g., the quadratic function) treats each voxel within a structure equally, whereas the EUD-based function aims to take into account of dose-volume effect when ranking the candidate plans. For a conventional approach to accomplish what an EUD method does, we partially “break” the implicit constraint that a structure is a uniform entity. The determination of the voxel-specific importance is dealt by heuristically relating them to the local doses. We have shown that voxel-based strategy is a most general ranking scheme, which can model both volumetric behavior and higher order factors. EUD and other dose-volume formalisms represent special cases of the

general framework. This work permits us to reformulate the EUD-based approach into the realm of conventional dose-based formalism without using biological parameters. With the new formalism, we can now design IMRT plans that would otherwise be unattainable.

Developments of other enabling technologies for the implementation of the novel adaptive prostate radiation therapy: To clinically implement ART, a number of “enabling” tools, such as deformable image registration, automated contour mapping from planning CT to the CBCT, and effective CBCT-based dose calculation algorithm, must be developed. In addition to tackle the dose optimization problem in ART by developing suitable dynamic control algorithms, we have also worked on establishing these enabling tools.

CBCT affords an effective means for us to examine the actual dose distribution to be delivered or already delivered to the patient on a routine basis. We have evaluated the accuracy of kV CBCT-based dose calculation and addressed some logistic issues related to its application for prostate IMRT. This work was presented in 2006 AAPM annual meeting in Orlando, FL, and a manuscript has been submitted to *Physics in Medicine and Biology* for publication.

Conventional deformable registration treats all image volume equally. Realized that some regions can be mapped between the moving and fixed images with higher confidence than others, we investigated a strategy of using *a priori* knowledge of the system to reduce the dimensionality of the deformable image registration problem and to speed up the registration calculation. Our approach consisted of two natural steps. First, a number of small cubic (0.5~1cm in size) control volumes are placed¹⁹, on the locally rigid regions of the moving image. Each control volume is mapped onto the moving image using a rigid transformation. In the second stage, the pre-determined correspondence serves as *a priori* information for the BSpline deformable registration calculation. This technique significantly reduces the search space and improves the convergence behavior of the gradient-based iterative optimization calculation. Along the same line, a contour mapping technique has been developed to map contours from simulation CT to daily CBCT images for adaptive therapy^{19, 20}.

Evaluation of potential impact of adaptive prostate radiation therapy: We have also investigated the potential impact of an integrated procedure for CBCT-based adaptive prostate IMRT treatment {Xing, 2006 #13}. Five prostate cases were selected for this study. For each patient, 3 gold markers were implanted into the prostate for target localization. A regular simulation CT was performed using a GE Discovery-ST scanner, and the prostate and seminal vesicle (SV) targets and the involved sensitive structures were outlined. 5-field IMRT plan was used for the actual patient treatment. Over the treatment

course, 3~5 CBCTs were acquired using Varian Trilogy after the patient was setup under the guidance of kV planar imaging (routine procedure). The CBCTs were registered to the sim-CT using a deformable model and the organ contours outlined on the sim-CT were auto-mapped to the CBCTs with tissue deformation taken into account. In addition to the actual plan for treatment, two plans were generated each time after a CBCT was acquired for off-line evaluation. These are: (i) the reconstructed delivered dose obtained using the beam parameters (fluence maps, beam configuration, and MUs) of the patient's treatment plan²¹; and (ii) the adaptive therapy plan with the fluence maps re-optimized to conform to the CBCT-derived, on-treatment patient anatomy and the dose delivery history reconstructed after each CBCT. In generating the 2nd plan, a gradient-search based dynamic algorithm was used to deal with the frequently updated inverse planning goals and patient geometric data. Comparison of the three types of plans was carried out to assess the improvement resulting from the adaptive approach. With the margins used in current prostate IMRT, we found that SV, bladder and rectum doses benefited most from the adaptive therapy. In addition, our study suggested that correcting the patient's daily setup just through the translation and rotation is often not enough and accounting for the organ deformation is important, especially if the target margin is to be reduced for dose escalation or for hypofractionated treatment. Adaptive therapy based on the volumetric on-board CT imaging and patient treatment history is an effective way to deliver highly conformal IMRT dose to prostate patients on a routine basis.

III. KEY RESEARCH ACCOMPLISHMENTS

- Established a theoretical relationship between spatially non-uniform penalty scheme and EUD formalism in inverse treatment planning.
- Developed method for incorporating a priori knowledge into deformable image registration.
- Evaluated the feasibility of using on-board cone-beam CT for on-treatment dose calculation.
- Developed novel closed-loop inverse planning strategies for adaptive prostate IMRT.
- Assessed the potential impact of adaptive radiation therapy for prostate cancer management.

IV. REPORTABLE OUTCOMES

The following is a list of publications resulted from the grant support in the last funding period.

Refereed publications:

1. de la Zerda A, Armbuster B, Xing L, Closed-Loop Control Algorithms for Planning Adaptive Radiation Therapy, *Physics in Medicine and Biology*, submitted, 2006.

2. Paquin D, Levy D, Schreibmann E., Xing L, Multistage image registration, *Mathematical Biosciences and Engineering* 3, 389-418, 2006. – figures featured in the cover of the issue of the journal.
3. Xing, L, Thorndyke B, Schreibmann E, Li T, Yang Y, Kim G., Luxton G, Koong, A, Overview of image guided radiation therapy (IGRT), *Medical Dosimetry* 31, 91-122, 2006.
4. Yang Y, Schreibmann E., Li T, Xing L, Evaluation of dosimetric accuracy of kV cone beam CT-based dose calculation, *Physics in Medicine and Biology*, submitted, 2006.
5. Chao M, Schreibamnn E, Li T, Wink N, Xing L, Automated contour mapping for 4D radiation therapy, *Medical Physics*, submitted, 2006.
6. Paquin D, Levy D, Xing L, Multistage deformable image registration, *IEEE Transactions on Medical Imaging*, submitted, 2006.
7. Paquin D, Levy D, Xing L, Hybrid of multistage and BSpline deformable image registration, *IEEE Transactions Image Processing*, submitted, 2006.

Conference abstract:

1. M Chao*, E Schreibmann, T. Li, L. Xing, Knowledge-Based Auto-Contouring in 4D Radiation Therapy, 2006 AAPM Annual Meeting, Orlando, FL.
2. Xing, L. and Spielman D, Functional and Molecular Imaging for Radiotherapy Guidance, 2006 AAPM Annual Meeting, Orlando, FL (invited talk).
3. B Armbruster, A de la Zerda, L Xing, A New 4D IMRT Algorithm and Its Performance Analysis, 2006 AAPM Annual Meeting, Orlando, FL.
4. B Armbruster, A de la Zerda, L Xing, A New Dose Optimization Algorithm for Adaptive Radiation Therapy, 2006 AAPM Annual Meeting, Orlando, FL.
5. M Chao,T Li, L Xing, Enhanced 4D CBCT Imaging for Slow-Rotating On-Board Imager, 2006 AAPM Annual Meeting, Orlando, FL.
2. Wang C., Yang Y, Schreibmann E., Li T., and Xing L: Evaluation of Kv CBCT-Based Dose Verification. 2006 AAPM Annual Meeting, Orlando, FL.
3. Yang Y., Levy D. and Xing L, Relationship of EDU-Based and Dose-Based Plan Optimization, 2006 AAPM Annual Meeting, Orlando, FL.
4. E Schreibmann*, B Thorndyke, L Xing, Intra- and Inter-Modality Registration of Four-Dimensional (4D) Images, 2006 AAPM Annual Meeting, Orlando, FL.
5. T Li, L Xing, P Munro, Y Yang, B Loo, A Koong, 4D Cone-Beam CT (CBCT) Using An On-Board Imager, 2006 AAPM Annual Meeting, Orlando, FL.
6. S Kamath, E Schreibmann, L Xing, Deformable Image Registration with Auto-Mapped Control Volumes, 2006 AAPM Annual Meeting, Orlando, FL.
7. A. de la Zerda, B. Armbruster and L. Xing, Inverse Planning for Adaptive Radiation Therapy Using Dynamic Algorithm, *International Journal of Radiation Oncology*Biology*Physics*, Volume 66, Issue 3, Supplement 1, 1 November 2006, Pages S123-S124.
8. W. Mao, T. Li, P. Munro, M. Chao and L. Xing, Individualizing 4D Cone-Beam CT (CBCT) Acquisition Protocol for External Beam Radiotherapy, *International Journal of Radiation Oncology*Biology*Physics*, Volume 66, Issue 3, Supplement 1, 1 November 2006, Pages S146-S147.
9. B. Cox, T. Ho, B. Thorndyke, T. Pawlicki, B. Loo, L. Xing, K. Goodman and A. Koong, Integrated Analysis of Pancreatic Tumor Motion Using Multiple Image-guided Modalities, *International Journal of Radiation Oncology*Biology*Physics*, Volume 66, Issue 3, Supplement 1, 1 November 2006, Pages S53-S54.
10. P. Lee, L. Xing, T. Pawlicki, P.T. Tran, A. Koong and K. Goodman, Image-Guided Radiation Therapy (RT) for Rectal Cancer Using Cone Beam CT (CBCT), *International Journal of Radiation Oncology*Biology*Physics*, Volume 66, Issue 3, Supplement 1, 1 November 2006, Page S276.
11. T. Li, E. Schreibmann, A. Koong, Q. Xu, R. Hamilton and L. Xing, Verification of Gated Radiation Therapy Using Pre-Treatment Four-Dimensional Cone-Beam CT, *International Journal of Radiation Oncology*Biology*Physics*, Volume 66, Issue 3, Supplement 1, 1 November 2006, Page S604.
12. L. Xing, A. de la Zerd, M. Cao, T. Li, B. Armbrush, Y. Yang, P. Lee, T. Pawlicki, S. Hancock and C. King, On-Board Volumetric CT-based Adaptive IMRT For Improved Prostate Cancer Treatment , *International Journal of Radiation Oncology Biology Physics* 66, 2006, S624-S625,

13. S. Kamath, E. Schreibmann, D. Levy, D. Paquin and L. Xing, Improving the Convergence and Computational Efficiency of Deformable Image Registration Calculation by Incorporating Prior Knowledge , *Int International Journal of Radiation Oncology Biology Physics* 66, 2006, S646-S647
14. D.C. Paquin, D. Levy and L. Xing, Multiscale Image Registration , *International Journal of Radiation Oncology Biology Physics* 66, 2006, S647
15. E. Schreibmann and L. Xing, Intra- And Inter-Modality Registration of Four-Dimensional (4D) Images *International Journal of Radiation Oncology Biology Physics* 66, 2006, , S648
16. M. Chao, E. Schreibmann, T. Li, A. Koong, K.A. Goodman and L. Xing, Automatic Contouring in 4D Radiation Therapy, *International Journal of Radiation Oncology Biology Physics* 66, 2006, S649.
17. M. Ding, L. Xing, W. Xiong, K. Stuhr and F. Newman, Validation of a 4D Monte Carlo Treatment Planning Tool Using an Image Interpolation Model • *International Journal of Radiation Oncology Biology Physics* 66, 2006, S655-S65.
18. E. Elder, E. Schreibmann, T. Li, T. Fox, L. Xing, J. Crocker and J. Landry, Registration of 4D CBCT and 4D CT for Extracranial Stereotactic Treatments • *International Journal of Radiation Oncology Biology Physics* 66, 2006, S651.
19. P. Peng, M. Chao, Q. Le, T. Li, A. Hsu, T.A. Pawlicki and L. Xing, Auto Contour Mapping in CBCT for ART Treatment Planning • *International Journal of Radiation Oncology Biology Physics* 66, 2006, S651-S652
20. I.R. Crocker, F. Tim, E. Elder, H. Shu, J. Landry, E. Schreibmann and L. Xing, Automated Segmentation of Cone Beam CT (CBCT) Datasets Using the Planning CT (PCT) as A-Priori Knowledge • *International Journal of Radiation Oncology Biology Physics* 66, 2006, S654.
21. B. Thorndyke, E. Schreibmann, A. Hsu, T. Fox and L. Xing, Four-Dimensional Registration of Respiratory Gated PET • *International Journal of Radiation Oncology Biology Physics* 66, 2006, S656-S657 .
22. C. Wang, Y. Yang and L. Xing, Evaluation of Cone Beam CT (CBCT)-Based Dose Calculation • *International Journal of Radiation Oncology Biology Physics* 66, 2006, S658-S659.
23. B. Armbruster, A. de la Zerda and L. Xing, Inverse Planning for 4D Intensity Modulated Radiation Therapy • *International Journal of Radiation Oncology Biology Physics* 66, 2006, S690-S691.
24. Yang Y, Levy D., Xing L, Voxel-based penalty scheme for inverse planning, invited talk in 2006 World Congress on Medical Physics and Bioengineering, Seoul, 2006.

V. CONCLUSIONS

Novel adaptive IMRT inverse planning formalism has been developed for the treatment of prostate cancer. A few important milestones have been achieved toward the general goal of the project. These include (i) Established a theoretical relationship between spatially non-uniform penalty scheme and EUD formalism in inverse treatment planning; (ii) Developed method for incorporating a priori knowledge into deformable image registration; (iii) Evaluated the feasibility of using on-board cone-beam CT for on-treatment dose calculation; (iv) Developed novel closed-loop inverse planning strategies for adaptive prostate IMRT; and (v) Assessed the potential impact of adaptive radiation therapy for prostate cancer management. Integration and further refinement of the above tools are in progress.

References:

1. Sonke, J. J.; Zijp, L.; Remeijer, P.; van Herk, M., Respiratory correlated cone beam CT. *Med Phys* 2005, 32, (4), 1176-86.
2. Li, T.; Xing, L.; McGuinness, C.; Munro, P.; Loo, B.; Koong, A., Four-dimensional cone-beam CT using an on-board imager. *Medical Physics* 2006, 33, 3825-3833.

3. Dietrich, L.; Jetter, S.; Tucking, T.; Nill, S.; Oelfke, U., Linac-integrated 4D cone beam CT: first experimental results. *Phys Med Biol* 2006, 51, (11), 2939-52.
4. Oldham, M.; Letourneau, D.; Watt, L.; Hugo, G.; Yan, D.; Lockman, D.; Kim, L. H.; Chen, P. Y.; Martinez, A.; Wong, J. W., Cone-beam-CT guided radiation therapy: A model for on-line application. *Radiother Oncol* 2005, 75, (3), 271 E1-8.
5. Pouliot, J.; Bani-Hashemi, A.; Chen, J.; Svatos, M.; Ghelmansarai, F.; Mitschke, M.; Aubin, M.; Xia, P.; Morin, O.; Bucci, K.; Roach, M., 3rd; Hernandez, P.; Zheng, Z.; Hristov, D.; Verhey, L., Low-dose megavoltage cone-beam CT for radiation therapy. *Int J Radiat Oncol Biol Phys* 2005, 61, (2), 552-60.
6. Yan, D.; Wong, J.; Vicini, F.; Michalski, J.; Pan, C.; Frazier, A.; Horwitz, E.; Martinez, A., Adaptive modification of treatment planning to minimize the deleterious effects of treatment setup errors. *Int J Radiat Oncol Biol Phys* 1997, 38, (1), 197-206.
7. Langen, K. M.; Meeks, S. L.; Poole, D. O.; Wagner, T. H.; Willoughby, T. R.; Kupelian, P. A.; Ruchala, K. J.; Haimerl, J.; Olivera, G. H., The use of megavoltage CT (MVCT) images for dose recomputations. *Phys Med Biol* 2005, 50, (18), 4259-76.
8. Yang, Y.; Schreibmann, E.; Li, T.; Xing, L., Dosimetric evaluation of kV cone-beam CT (CBCT) based dose calculation. *Medical Physics* 2006, submitted.
9. Trofimov, A.; Rietzel, E.; Lu, H. M.; Martin, B.; Jiang, S.; Chen, G. T.; Bortfeld, T., Temporo-spatial IMRT optimization: concepts, implementation and initial results. *Phys Med Biol* 2005, 50, (12), 2779-98.
10. Wu, Q.; Liang, J.; Yan, D., Application of dose compensation in image-guided radiotherapy of prostate cancer. *Phys Med Biol* 2006, 51, (6), 1405-19.
11. Mohan, R.; Zhang, X.; Wang, H.; Kang, Y.; Wang, X.; Liu, H.; Ang, K. K.; Kuban, D.; Dong, L., Use of deformed intensity distributions for on-line modification of image-guided IMRT to account for interfractional anatomic changes. *Int J Radiat Oncol Biol Phys* 2005, 61, (4), 1258-66.
12. de la Zerda, A.; Ambrush, B.; Xing, L., Inverse planning for adaptive radiation therapy using dynamic programming. In *Annual Meeting of ASTRO*, Philadelphia, PA, 2006; Vol. 32, p 2163.
13. Olivera, G. H.; Mackie, T. R.; Ruchala, K.; Lu, W.; Kapatoes, J., Adaptive radiation therapy (art) strategies using helical tomotherapy. In *Image-Guided IMRT*, Bortfeld, T.; Schmidt-Ullrich, R.; De Deve, W.; Wazer, D. E., Eds. Springer-Verlag: Berlin, 2006; pp 235-246.
14. Ambrush, B.; de la Zerda, A.; Levy, D.; Xing, L., Inverse planning for 4D radiation therapy. In *Annual Meeting of AAPM*, AAPM: Orlando, FL, 2006; Vol. 32, p 2163.
15. Widrow, B.; Stearns, S., *Adaptive Signal Processing*. Prentice Hall: Upper Saddle River, NJ, 1985.
16. Widrow, B.; Walach, E., *Adaptive Inverse Control*. Prentice Hall: Upper Saddle River, NJ, 1995.
17. Yang, Y.; Xing, L., Optimization of radiation dose-time-fractionation scheme with consideration of tumor specific biology. *Medical Physics* 2005, 32, (), 3666-3677.
18. Li, T.; Schreibmann, E.; Yang, Y.; Xing, L., Motion correction for improved target localization with on-board cone-beam computed tomography. *Phys Med Biol* 2006, 51, (2), 253-67.
19. Cao, M.; Schreibmann, E.; Paquin, D.; Levy, D.; Xing, L., Mapping of organ lung between different phases in 4D CT images. In *Annual Meeting of AAPM*, AAPM: Orlando, FL, 2006; Vol. 32, p 2163.
20. Cao, M.; Schreibmann, E.; Paquin, D.; Levy, D.; Xing, L., Automated segmentation for adaptive radiation therapy. In *Annual Meeting of AAPM*, AAPM: Orlando, FL, 2006; Vol. 32, p 2163.
21. Xing, L.; de la Zerd, A.; Chao, C.; Li, T.; Armbrush, A.; Yang, Y.; Lee, P.; Pawlicki, T.; Hancock, S.; King, C., On-Board Volumetric CT-based Adaptive IMRT For Improved Prostate Cancer Treatment *International Journal of Radiation Oncology Biology Physics*, 2006, 66, S624-S625

Appendices -- manuscripts published/submitted

1. de la Zerda A, Armbruster B, Xing L, Closed-Loop Control Algorithms for Planning Adaptive Radiation Therapy, *Physics in Medicine and Biology*, submitted, 2006.
2. Paquin D, Levy D, Schreibmann E., Xing L, Multistage image registration, *Mathematical Biosciences and Engineering* 3, 389-418, 2006.
3. Xing, L, Thorndike B, Schreibmann E, Li T, Yang Y, Kim G., Luxton G, Koong, A, Overview of image guided radiation therapy (IGRT), *Medical Dosimetry* 31, 91-122, 2006.
4. Yang Y, Schreibmann E., Li T, Xing L, Evaluation of dosimetric accuracy of kV cone beam CT-based dose calculation, *Physics in Medicine and Biology*, submitted, 2006.
5. Chao M, Schreibamnn E, Li T, Wink N, Xing L, Automated contour mapping for 4D radiation therapy, *Medical Physics*, submitted, 2006.
6. Paquin D, Levy D, Xing L, Multistage deformable image registration, *IEEE Transactions on Medical Imaging*, submitted, 2006.
7. Paquin D, Levy D, Xing L, Hybrid of multistage and BSpline deformable image registration, *IEEE Transactions Image Processing*, submitted, 2006.

Closed-Loop Control Algorithms

for Planning Adaptive Radiation Therapy

Adam de la Zerda, Benjamin Armbruster, and Lei Xing^a

Department of Radiation Oncology, Stanford University, Stanford, CA 94305-5847

Short title: Inverse Planning for Adaptive Therapy

^{a)} Corresponding author:
Department of Radiation Oncology
Stanford University School of Medicine
Clinical Cancer Center
875 Blake Wilbur Drive
Stanford, CA 94305-5847
Telephone: (650) 498-7896
Fax: (650) 498-4015
Email: lei@reyes.stanford.edu

Submitted to: *Physics in Medicine and Biology*

Abstract

Current radiation therapy (RT) does not adapt to inter-fraction organ movement and dosimetric errors caused by inaccurate setup or organ deformation during a course of treatment. The emergence of on-board cone beam CT (CBCT) affords an effective means to obtain the patient's geometric model just before treatment and recompute on a routine basis the dose to be delivered or actually delivered to the patient. This makes it possible to adaptively correct for dosimetric errors in the previous fractions by modifying the treatment plan. However, before this new scheme of RT can happen clinically, an inverse planning algorithm capable of taking into account the dose delivery history and the patient's geometric model must be in place. In this paper we devise dynamic closed-loop control algorithms for adaptive therapy (ART) and demonstrate their utility with data from phantom and clinical cases. To meet the needs of different clinical applications, we study two classes of algorithms: those *Adapting to Changing Geometry* and those *Adapting to Geometry and Delivered Dose*. The former class takes into account organ deformations found just before treatment. The latter class optimizes the dose distribution accumulated over the entire course of treatment by adapting at each fraction not only to the information just before treatment about organ deformations but also to the dose delivery history. We showcase two algorithms in the class of those *Adapting to Geometry and Delivered Dose*. We study the feasibility and utility of the algorithms using phantom and clinical cases. A comparison with conventional approaches indicates that ART optimization may significantly improve the current practice.

Keywords: adaptive therapy, IGRT, closed-loop control, inverse planning, organ motion, IMRT

1. Introduction

Current IMRT treatment plan optimization and dose delivery are two decoupled steps. In each fraction, the patient geometry is hardly the same as in the pre-treatment CT simulation. A commonly used method to take the uncertainty into account is to add a safety margin, whose size is based on population statistics, to the target and sensitive structures. This significantly compromises the success of radiation therapy.

Recently, cone-beam computed tomography (CBCT) integrated with a medical linear accelerator has become available and promises to improve the situation. CBCT provides a valuable 3D (or even possibly 4D(Dietrich et al., 2006; Li et al., 2006; Sonke et al., 2005)) geometric model of the patient in the treatment position and allows for verification of the delivered dose distribution. This not only affords an opportunity for online correction of patient setup error and inter-fraction rigid motion(Oldham et al., 2005; Pouliot et al., 2005), but also allows dose recalculation (Langen et al., 2005; Yan et al., 1997; Yang et al., 2006), and adaptive radiation therapy (ART)(Langen et al., 2005; Yan et al., 1997; Yang et al., 2006), which uses the volumetric information to adjust the treatment plan each fraction to the updated patient anatomy and positioning. A significant promise of ART is the optimal compensation of uncertainties including organ deformation and inter-fraction organ motion as well as dosimetric errors incurred in previous fractions(de la Zerda et al., 2006; Mohan et al., 2005; Olivera et al., 2006; Trofimov et al., 2005; Wu et al., 2006). To realize ART clinically and maximally exploit the potential of this new form of image guided radiation therapy (IGRT), a robust inverse planning strategy for ART must be in place.

The purpose of this work is to develop dynamic closed-loop control strategies for ART inverse planning and demonstrate their utility with data from phantom and clinical cases.

Closed-loop control algorithms are a general tool for dealing with time-dependent systems (Widrow and Stearns, 1985; Widrow and Walach, 1995) and are used to solve a variety of problems including automobile cruise control, supply chain optimization, computer chip design, and spaceship navigation. The algorithms in all these applications share the same basic closed-loop control framework of repeated re-evaluation and re-planning. ART is a natural application for closed-loop control because CBCT provides frequently updated system information. To meet different clinical requirements, we investigate a few variants of these algorithms.

The next section starts with a brief review of the foundations of closed-loop control. Section B continues by presenting a formalism for ART planning, the intensity-modulated radiation therapy (IMRT) approach, and our new ART algorithms. In section C, we demonstrate the significance of our new approach with phantom and patient case studies. Section D takes a broader perspective to discuss this work and directions for future research. We conclude in section E.

2. Methods and Materials

This section starts with an overview of closed-loop control and then later describes various closed-loop control algorithms for ART before summarizing these algorithms in Table 2.

2.1 Foundations of Closed-Loop Control

The use of feedback makes closed-loop control unique(Widrow and Stearns, 1985; Widrow and Walach, 1995). In open-loop control, Fig. 1a, some *input* is fed to the *controller* which then decides on an *action*. The controller then performs that action on the *system* (the system being the physical or biological system we want to control)

effectively putting the system into a new state (which hopefully is “better” in some sense). We could simulate an open-loop control context by connecting the output of the algorithm implemented by the controller to a simulation of the system: the input is fed to the algorithm which feeds an action to the system simulator which outputs the new state of the system.

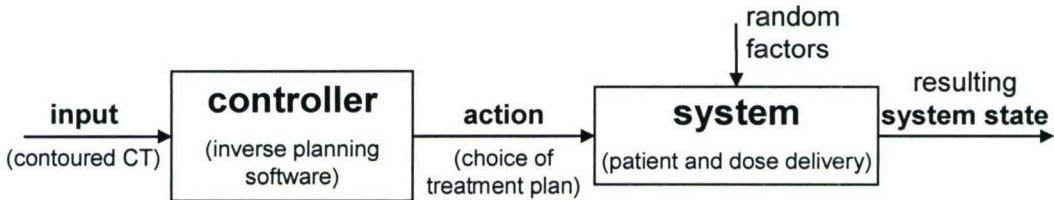


Figure 1a: open-loop control (conventional treatment planning)

Conventional inverse planning is a form of open-loop control. The *input* is a set of contoured CT images or the geometric model of the patient. The *controller* is the inverse planning software, and it outputs a treatment plan (the *action*). Finally, a simulator of the system (modeling the execution of the plan) uses the treatment plan and some random factors to output the resulting system state.

In closed-loop control, the controller is not run just once but repeatedly, each time receiving the current state of the system as its input (Widrow and Stearns, 1985; Widrow and Walach, 1995). The term feedback is used because the effect of the current action (the new state of the system) is used by the controller to plan the next action. Suppose we have N time periods (treatment fractions), and we let β_i be the action in period i (treatment plan for fraction i) and ψ_i the resulting system state at the end of period i . Then a closed-loop controller is a function producing an action for the current period

from the state at the end of the last period $\beta_i = \xi_i(\psi_{i-1})$.

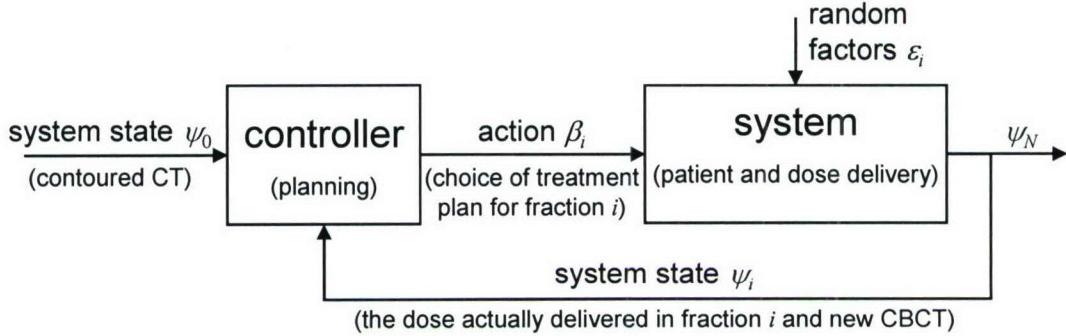


Figure 1b: the closed-loop control framework

We propose to plan ART using closed-loop control, Fig. 1b, where the controller is the inverse planning software and the state of the system after fraction i , $\psi_i = (\psi_i^{\text{geometry}}, \psi_i^{\text{cumdose}})$, has two parts: the patient's geometric model derived from contoured CBCT images, ψ_i^{geometry} , and the cumulative dose, ψ_i^{cumdose} , delivered in fractions 1 through i . (Clearly, $\psi_0^{\text{cumdose}} = 0$.) Unlike the open-loop control in conventional RT, in ART the controller outputs a plan (the action) for only the current treatment fraction. Table 1 summarizes the relationship between control theory and treatment planning.

Closed-loop control algorithms (which we propose for ART) are richer than static open-loop control algorithms (found for example in conventional RT). A static open-loop controller determines a treatment plan, $(\beta_1, \dots, \beta_N)$, and therefore the final state of the system ψ_N based only on the initial state of the system ψ_0 , whereas dynamic closed-loop algorithms reevaluate the state each time step and based on that, decide their

current decision. Because of that, the treatment plan, $(\beta_1, \dots, \beta_N)$, depends not only on the initial state ψ_0 , but also on the progress made during treatment, $(\psi_0, \dots, \psi_{N-1})$. Hence closed-loop control algorithms are more likely to produce better results than static open-loop control algorithms for time-dependent systems.

Treatment planning is a critical step to realizing the potential of ART in clinical practice. Current IMRT inverse planning is designed for a conventional fractionated treatment scheme (Alber and Nusslin, 1999; Bedford and Webb, 2003; Bortfeld, 1999; Censor, 2003; Hristov and Fallone, 1997; Oelfke and Bortfeld, 1999; Rowbottom et al., 2001; Shepard et al., 2002; Soderstrom et al., 1993; Thieke et al., 2003; Webb et al., 1998; Xiao et al., 2000; Xing et al., 2005; Yang and Xing, 2004b) and is incapable of planning an adaptive treatment with consideration of the dose delivery history. The purpose of this work is to establish a dynamic control framework for adaptive radiation therapy. This planning framework allows us to individualize radiation therapy for each patient by taking advantage of newly available volumetric imaging information.

Closed-Loop Control	Radiation Therapy
time period (i)	fraction
input (ψ_0)	contoured CT images, $(\psi_0^{\text{geometry}})$
action in period i , (β_i)	treatment plan for fraction i
system state after period i , (ψ_i)	contoured CBCT images, $(\psi_i^{\text{geometry}})$, and cumulative delivered dose, $(\psi_i^{\text{cumdose}})$
controller	RT inverse planning software
system	patient geometry and treatment delivery

random factors/noise in period i , (ε_i)	setup errors, delivery errors, and organ deformations in fraction i
--	---

Table 1: Relationship between control theory and treatment planning

2.2 Formalization of ART plan optimization

Consider plans, $(\beta_1, \dots, \beta_N)$, for N fractions. Suppose at voxel v , the importance factor is $\alpha(v)$ (Alber and Nusslin, 1999; Corletto et al., 2003; Li and Yin, 2000; Xing et al., 1999) and the prescribed dose is $D^{\text{prescribed}}(v)$. We define a dose delivery function $D(\cdot)$, such that under plan β and delivery error ε the cumulative dose after fraction i is $\psi_i^{\text{cumdose}}(v) = \psi_{i-1}^{\text{cumdose}}(v) + D(v; \beta, \varepsilon, \psi_{i-1}^{\text{geometry}})$ at voxel v . Without delivery errors (i.e., without any setup errors or deformations of organs), $\varepsilon = 0$. Let B be the set of deliverable (feasible) plans. Our ultimate goal is then to find the feasible plan $(\beta_1, \dots, \beta_N)$ optimizing

$$\begin{aligned} & \min_{\beta_1 \in B, \dots, \beta_N \in B} \sum_v \alpha(v) \left(D^{\text{prescribed}}(v) - \psi_N^{\text{cumdose}}(v) \right)^2 \\ &= \min_{\beta_1 \in B, \dots, \beta_N \in B} \sum_v \alpha(v) \left(D^{\text{prescribed}}(v) - \sum_{i=1}^N D(v, \beta_i, \varepsilon_i, \psi_{i-1}^{\text{geometry}}) \right)^2. \end{aligned} \quad (1)$$

The objective function is the weighted quadratic deviation of the cumulative delivered dose from the prescribed dose. There is no way to find in advance the optimal solution to this problem because the actual delivery error in fraction i , ε_i , is unknown when we must decide on the plan β_i . While the simplest idea is to optimize

$$\min_{\beta_1 \in B, \dots, \beta_N \in B} \sum_v \alpha(v) \left(D^{\text{prescribed}}(v) - \sum_{i=1}^N D(v, \beta_i, 0, \psi_0^{\text{geometry}}) \right)^2,$$

we suggest in the following sections some more sophisticated approximations to

problem (1).

2.3 Baseline Algorithm 1: Planning with Population-Based Margins

A common approach adds large margins to structures to compensate for delivery errors and uses the same plan every fraction, $\beta_1 = \dots = \beta_N$ (i.e., a static plan). These margins lead to modified dose prescriptions, $\tilde{D}^{\text{prescribed}}(\cdot)$, and importance factors, $\tilde{\alpha}(\cdot)$. This yields the new optimization problem:

$$\beta_1 = \dots = \beta_N = \arg \min_{\beta \in B} \sum_v \tilde{\alpha}(v) \left(\frac{\tilde{D}^{\text{prescribed}}(v)}{N} - D(v, \beta, 0, \psi_0^{\text{geometry}}) \right)^2.$$

While this approximation is simple, it exposes much healthy tissue to high dose.

2.4 Baseline Algorithm 2: Adapting to Changing Geometry

Assume that a CBCT imaging is taken prior to every treatment fraction. Based on the current anatomy information, a new plan for the current fraction is formed. This means that we consider $\psi_{i-1}^{\text{geometry}}$ when determining β_i . In particular we choose for fraction i the feasible plan that minimizes the weighted quadratic deviation from the initial prescribed dose: for all fractions i ,

$$\beta_i = \arg \min_{\beta \in B} \sum_v \alpha(v) \left(\frac{D^{\text{prescribed}}(v)}{N} - D(v, \beta, 0, \psi_{i-1}^{\text{geometry}}) \right)^2.$$

In other words, the algorithm tries to deliver the same daily prescribed dose, $D^{\text{prescribed}}/N$, in each fraction.

2.5 Adapting to Geometry and Delivered Dose

We can do better than in the previous section, B.4, by taking into account not only the

up-to-date CBCT imaging information $\psi_{i-1}^{\text{geometry}}$, but also other factors such as the cumulative delivered dose $\psi_{i-1}^{\text{cumdose}}$. In this work we present two algorithms that take into account both $\psi_{i-1}^{\text{geometry}}$ and $\psi_{i-1}^{\text{cumdose}}$ when choosing a feasible plan for fraction i . In general this adaptive dose goal may depend on the dose delivery history, $\psi_{i-1}^{\text{cumdose}}$, and predictions of future geometry. The two algorithms we investigate differ in how they choose this adaptive dose goal. The first compensates every fraction for delivery errors in previous fractions, and the second incorporates predictions of the future patient geometry.

2.5.1 Immediately Correcting Algorithm (ICA)

This algorithm, Fig. 2, takes into account the dose delivery history as well as the anatomy model derived from daily CBCT images. No prediction is attempted about the patient model in the subsequent fractions. This algorithm adjusts the originally prescribed dose to completely compensate voxels which were overdosed (or underdosed) in previous fractions by decreasing (or increasing) the dose goal at those voxels. Specifically, it adds in fraction i the accumulated error from the previous

fractions, $\frac{i-1}{N} D^{\text{prescribed}} - \psi_{i-1}^{\text{cumdose}}$, to the original daily prescription dose, $D^{\text{prescribed}}/N$,

resulting in an adaptive dose goal of

$$D_i^{\text{goal}} = \frac{D^{\text{prescribed}}}{N} + \left(\frac{i-1}{N} D^{\text{prescribed}} - \psi_{i-1}^{\text{cumdose}} \right) = \frac{i}{N} D^{\text{prescribed}} - \psi_{i-1}^{\text{cumdose}}.$$

Note that the adaptive dose goal can be accurately calculated as long as the previous delivered doses, $\psi_{i-1}^{\text{cumdose}}$, are known. While this is currently not the case, the information is becoming increasingly available. This algorithm chooses the feasible plan $\beta_i \in B$ for fraction i that minimizes the weighted quadratic deviation of the planned

dose, $D(\cdot, \beta_i, 0, \psi_{i-1}^{\text{geometry}})$ from the adaptive dose goal, D_i^{goal} :

$$\beta_i = \arg \min_{\beta \in B} \sum_v \alpha(v) (D_i^{\text{goal}}(v) - D(v, \beta, 0, \psi_{i-1}^{\text{geometry}}))^2 \quad (2)$$

$$= \arg \min_{\beta \in B} \sum_v \alpha(v) \left(\frac{i}{N} D^{\text{prescribed}}(v) - \psi_{i-1}^{\text{cumdose}}(v) - D(v, \beta, 0, \psi_{i-1}^{\text{geometry}}) \right)^2.$$

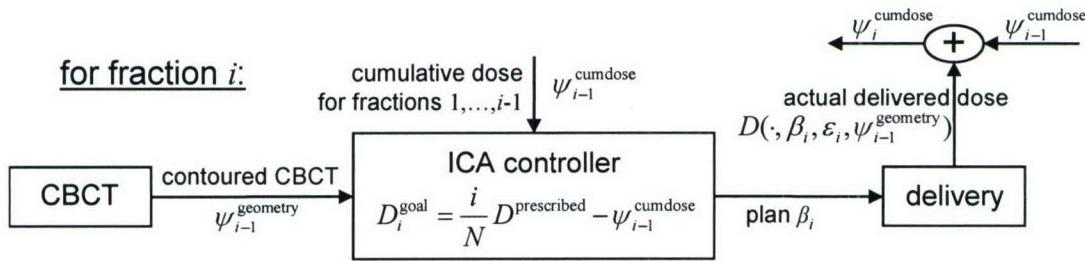


Figure 2: Block Diagram for ICA

2.5.2 Prudent Correcting Algorithm (PCA)

Consider a treatment course of N fractions. Let i be the current fraction index and d be the number of fractions for which we forecast the patient's anatomy and position (e.g., $d = 2$ means predicting two days ahead of the current fraction). So to deliver the current fraction we find the optimal plan not only for this fraction but also for the subsequent d fractions. Note, that d may be a function of the current fraction i . For instance, taking $d(i) = N - i$ means that we predict the anatomy and organ locations for the rest of the treatment course. Optimizing only for the current fraction ignores our (albeit imperfect) knowledge of the future and therefore misses the opportunity to compensate the dose in future fractions. Note that this does not prevent one from updating the system when the next set of information becomes available at a certain future fraction. To determine the plan chosen for fraction i we optimize

$$\min_{\beta'_i \in B, \dots, \beta'_{i+d} \in B} \sum_v \alpha(v) \left(\frac{i+d}{N} D^{\text{prescribed}}(v) - \psi_{i-1}^{\text{cumdose}}(v) - \sum_{j=i}^{j=i+d} D(v, \beta'_j, 0, \psi_j^{\text{predicted-geometry}}) \right)^2 \quad (3)$$

where $\psi_i^{\text{predicted-geometry}}, \dots, \psi_{i+d}^{\text{predicted-geometry}}$ are the predicted anatomy locations for fractions i through $i+d$. The optimizers of problem (3) are a sequence of plans, $\beta'_i, \dots, \beta'_{i+d}$, the first of which is the plan, β_i , that we will implement in fraction i . Clinically, this scheme is useful when dealing with situations of tumor shrinkage after a replanning CT/CBCT is done (see the discussion section for more details).

We consider a simple prediction model in which the anatomy remains unchanged for the next d fractions to illustrate the approach (in practice one may use more sophisticated prediction models). That is, we assume $\psi_j^{\text{predicted-geometry}} = \psi_{i-1}^{\text{geometry}}$ for fractions i through $i+d$. For this prediction model, it can be shown that (3) reduces to

$$\beta_i = \arg \min_{\beta \in B} \sum_v \alpha(v) \left(\frac{i+d}{N} D_i^{\text{prescribed}}(v) - \psi_{i-1}^{\text{cumdose}}(v) - (1+d) D(v, \beta, 0, \psi_{i-1}^{\text{geometry}}) \right)^2.$$

We can also give this model the form of problem (2) by defining

$$D_i^{\text{goal}} = \frac{D^{\text{prescribed}}}{N} + \frac{1}{d+1} \left(\frac{i-1}{N} D^{\text{prescribed}} - \psi_{i-1}^{\text{cumdose}} \right). \quad (4)$$

If $d = 0$, then this algorithm coincides with the previous ICA algorithm. In equation (4) we see that this algorithm differs from the ICA algorithm because the correction to the accumulated error is divided among the subsequent $d+1$ fractions of therapy to achieve better uniformity and robustness of therapy. Alternative schedules for compensating the accumulated error that accommodate specific clinical considerations should be easily implementable.

Algorithm	Objective functions for different schemes of dose optimization
Perfect Foresight	$\min_{\beta_1 \in B, \dots, \beta_N \in B} \sum_v \alpha(v) \left(D^{\text{prescribed}}(v) - \sum_{i=1}^N D(v, \beta_i, \epsilon_i, \psi_{i-1}^{\text{geometry}}) \right)^2$ <p>minimize the difference between the prescription and the delivered dose accumulated over all treatment fractions</p>
Baseline 1: Population-Based Margins	$\beta_1 = \dots = \beta_N = \arg \min_{\beta \in B} \sum_v \tilde{\alpha}(v) \left(\frac{\tilde{D}^{\text{prescribed}}(v)}{N} - D(v, \beta, 0, \psi_0^{\text{geometry}}) \right)^2$ <p>Add margins to the prescription and then minimize the difference between the prescription and the daily planned dose.</p>
Baseline 2: Adapting to Changing Geometry	$\beta_i = \arg \min_{\beta \in B} \sum_v \alpha(v) \left(\frac{D^{\text{prescribed}}(v)}{N} - D(v, \beta, 0, \psi_{i-1}^{\text{geometry}}) \right)^2$ <p>Update the patient's geometric model every fraction using CBCT and plan a dose for that geometry that minimizes the difference to the daily prescription.</p>
Adapting to Geometry and Delivered Dose	$\beta_i = \arg \min_{\beta \in B} \sum_v \alpha(v) \left(\frac{i}{N} D^{\text{prescribed}}(v) - \psi_{i-1}^{\text{cumdose}}(v) - D(v, \beta, 0, \psi_{i-1}^{\text{geometry}}) \right)^2$ <p>Update the patient's geometric model every fraction using CBCT and plan a dose for that geometry that minimizes the difference to the prescribed dose plus the accumulated error.</p>
	$\beta_i = \arg \min_{\beta \in B} \sum_v \alpha(v) \left(\frac{i+d}{N} D_i^{\text{prescribed}}(v) - \psi_{i-1}^{\text{cumdose}}(v) - (1+d) D(v, \beta, 0, \psi_{i-1}^{\text{geometry}}) \right)^2$ <p>Update the patient's geometric model every fraction using CBCT and using that geometry plan a dose for remaining ($d+1$) fractions that minimizes the difference to the prescribed dose plus the accumulated error.</p>

Table 2: Summary of various radiation therapy algorithms

2.6 Evaluation and Case Study

We developed an in-house inverse planning platform to evaluate this closed-loop control framework and the novel *Adapting to Geometry and Delivered Dose* algorithms. This platform implements various RT planning strategies by optimizing fluence maps using a commercial nonlinear optimization code SNOPT (Gill et al., 2005). The platform also

simulates the effectiveness of these plans by implementing the dose delivery function $D(\cdot)$.

Using this platform we compare the following algorithms: *Adapting to Changing Geometry* (our baseline), *Immediately Correcting Algorithm* (ICA), *Prudent Correcting Algorithm* (PCA), and *Perfect Foresight*. Results of the Population-Based Margins Planning Algorithm are not shown since the selection of margin size varies with institution and, in general, it performs much worse than our baseline. *Perfect Foresight* is an algorithm that accurately predicts future errors and motions (see Table 2). The performance of this algorithm may not actually be attainable, but it sets a theoretical upper bound on the performance of any RT planning algorithm.

We compared the algorithms on a cubic phantom case and a prostate patient case. The phantom, Fig. 3, is 20cm in size and at its center has a C-shaped target enclosing a round sensitive structure. In both cases, we asked the algorithms to generate a plan with 15 fractions and using 5 equally spaced beam directions. The tumor, sensitive structures, and normal tissue were assigned importance factors of $\alpha = 10$, $\alpha = 3$, and $\alpha = 0.7$ respectively. We simulated setup errors by introducing independently for every fraction a random translation chosen uniformly at random from [-1, 1]cm and a random rotation chosen uniformly at random from [-2 $^{\circ}$, 2 $^{\circ}$]. In section C we compare the results visually in terms of the cumulative DVH and numerically for each organ in terms of the average dose and standard deviation of the dose.

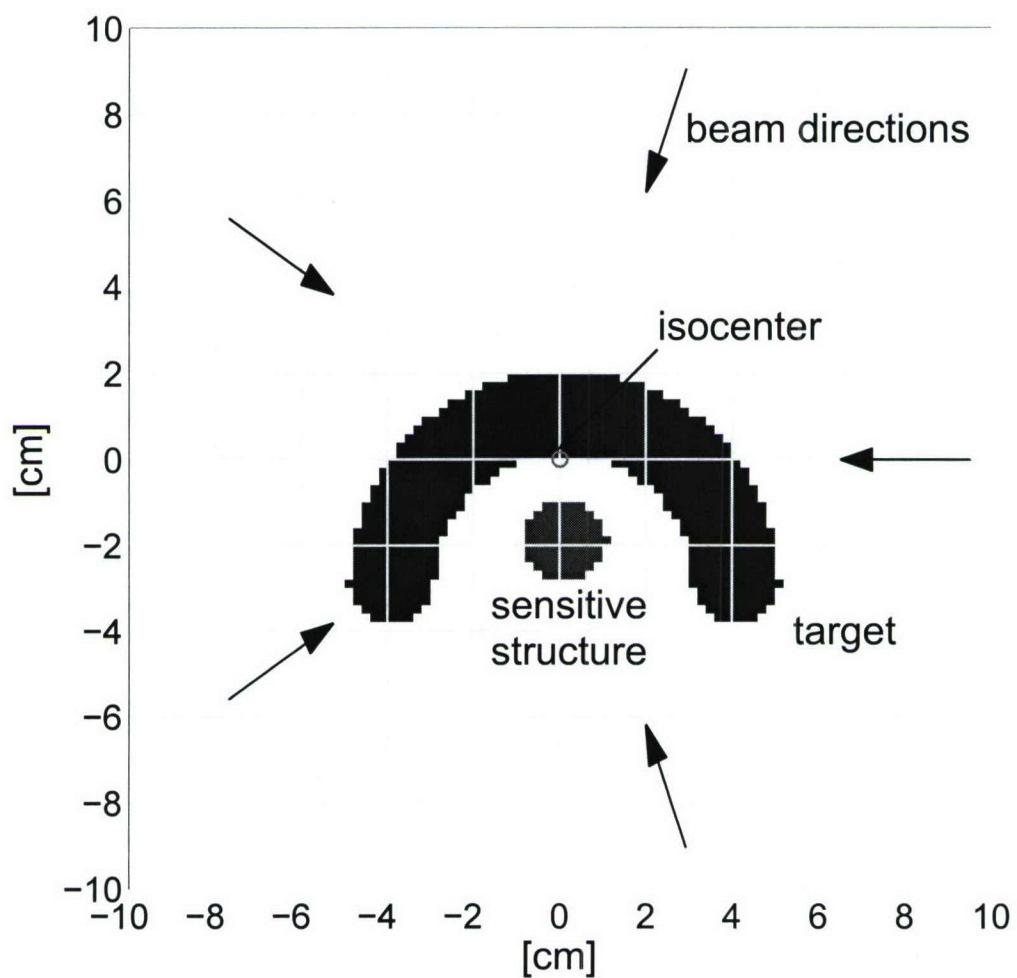


Figure 3: Phantom Anatomy

3. Results

3.1 Phantom study

The results are shown graphically in Fig. 4 and statistically in Table 3. Each row of table 3 compares an algorithm to our baseline, *Adapting to Changing Geometry*.

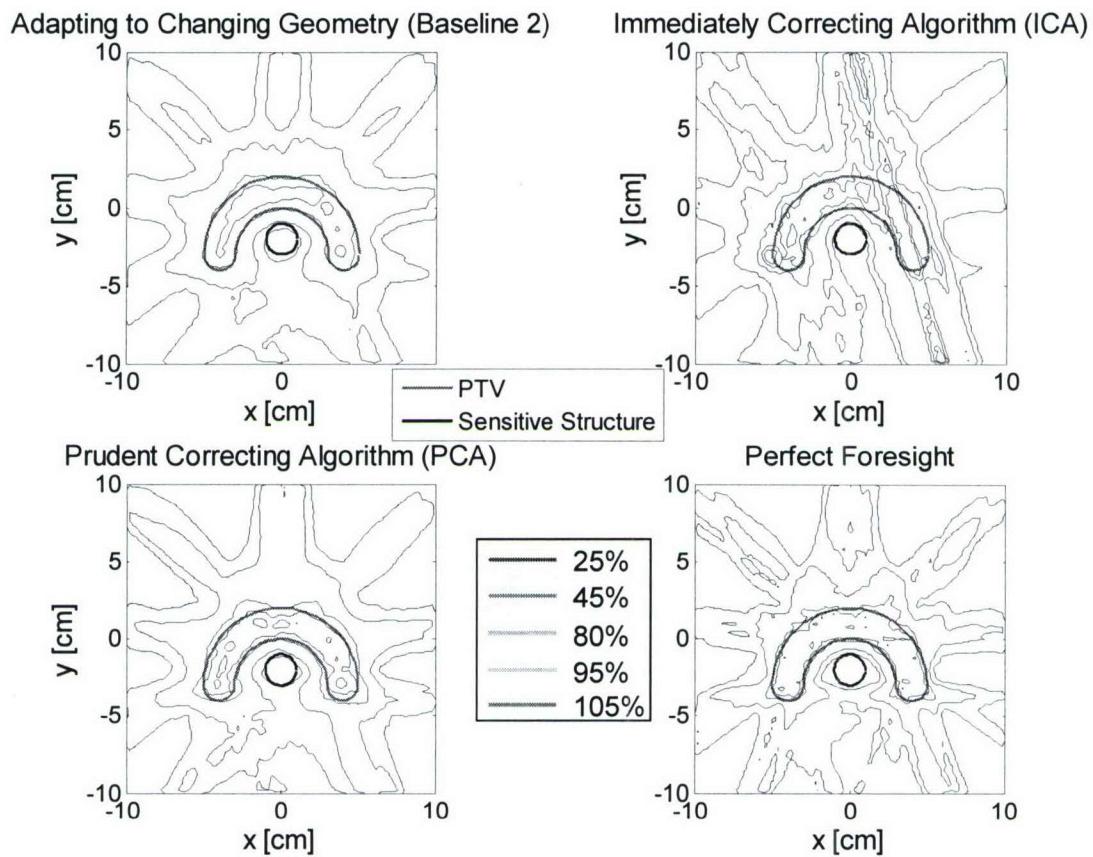


Figure 4a: Cumulative dose delivered to phantom as percent of prescription

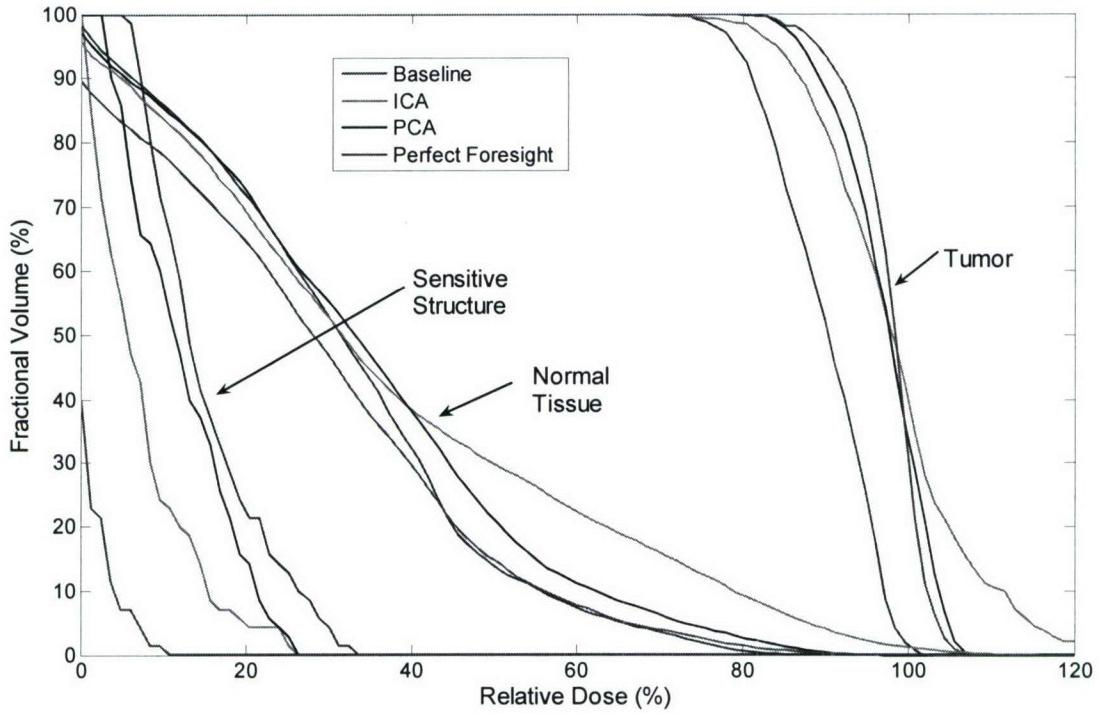


Figure 4b: DVHs of cumulative dose delivered to phantom as percent of prescription.

Algorithm	Tumor		Sensitive Structure		Normal Tissue	
	Avg. Dose	1σ	Avg. Dose	1σ	Avg. Dose	1σ
ICA	+9%	+48%	-48%	-12%	+18%	+47%
PCA	+8%	-13%	-18%	-9%	+7%	+14%
Perfect Foresight	+9%	-30%	-89%	-65%	-8%	+13%

Table 3: Cumulative dose delivered to phantom for various algorithms relative to the

baseline algorithm, *Adapting to Changing Geometry*.

Fig. 4a shows three benefits the *Adapting to Geometry* and *Delivered Dose* algorithms (and of course the *Perfect Foresight* algorithm) have over the baseline *Adapting to Changing Geometry*. For the first benefit, notice how the 25% contour in our baseline is shifted slightly down from the sensitive structure. All the other algorithms do not show this shift. This shift is most likely due to some setup error (that isn't corrected for in our baseline). The second benefit (also seen in the DVHs and Table 3) is dose escalation to

the tumor while keeping the sensitive structure close to 25% (unlike the baseline, these algorithms achieve tumor doses close to 100%). The third benefit is steeper gradients around the tumor (because dose correcting algorithms won't let errors accumulate).

Aside from these benefits, the dose distributions in Fig. 4a also show how much worse ICA is than PCA. As discussed above, ICA is able to escalate dose to the tumor while keeping the dose to the sensitive structure low, but along the way it overdoses normal tissue and decreases dose uniformity (see Table 3). ICA may have fractions with higher doses than PCA since it tries to completely correct for dose delivery errors in the next fraction. We believe this makes ICA less robust than PCA because when combined with organ deformation or setup error these higher doses lead to bigger errors. The differing performance of the two *Adapting to Geometry and Delivered Dose* algorithms show that the algorithmic details matter.

Looking at the DVH in Fig. 4b, we see that the *Perfect Foresight* algorithm achieves significantly better results than the adaptive algorithms (e.g., PCA). This is explained by the obvious fact that the *Perfect Foresight* algorithm has the fundamental advantage of accurately predicting the future. This gap between the adaptive algorithms and the *Perfect Foresight* algorithm represents the potential of closed-loop control algorithms as they become more sophisticated in terms of future predictions and therefore become closer and closer to the *Perfect Foresight* algorithm. Especially the dose to the sensitive structure and to the normal tissue is not as low as in the *Perfect Foresight* result. In addition, the tumor dose uniformity could improve a bit.

3.2 Prostate Study

The resulting dose distributions and their DVHs are in figures 5a and 5b. Each row of

table 4 compares an algorithm to our baseline, *Adapting to Changing Geometry*.

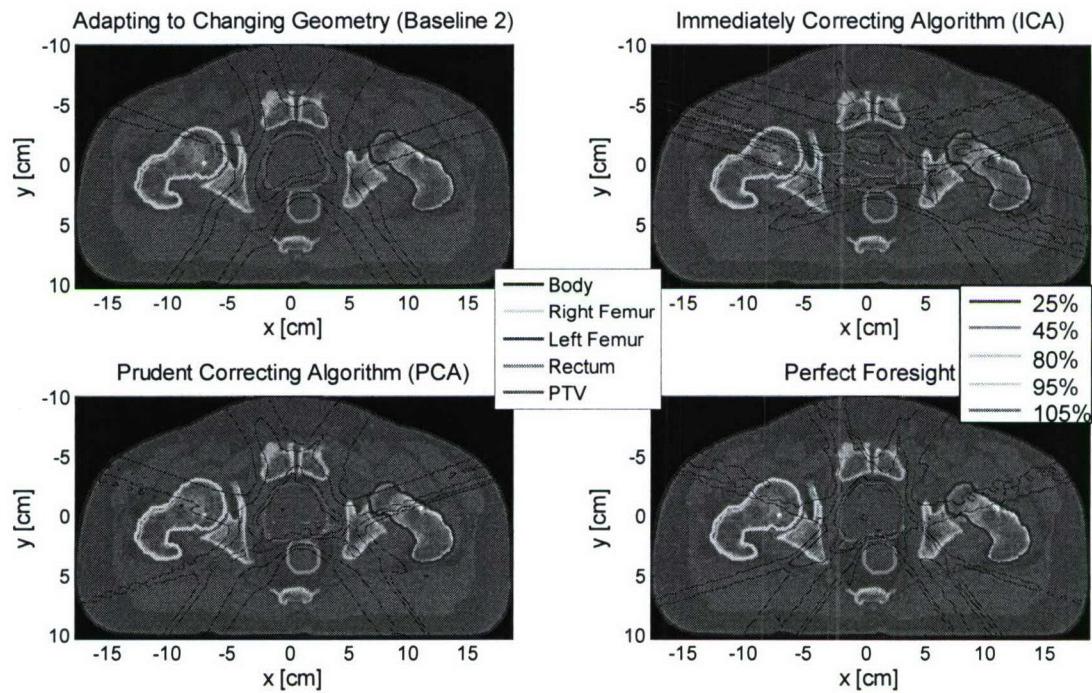


Figure 5a: Cumulative dose delivered in prostate study as percent of prescription

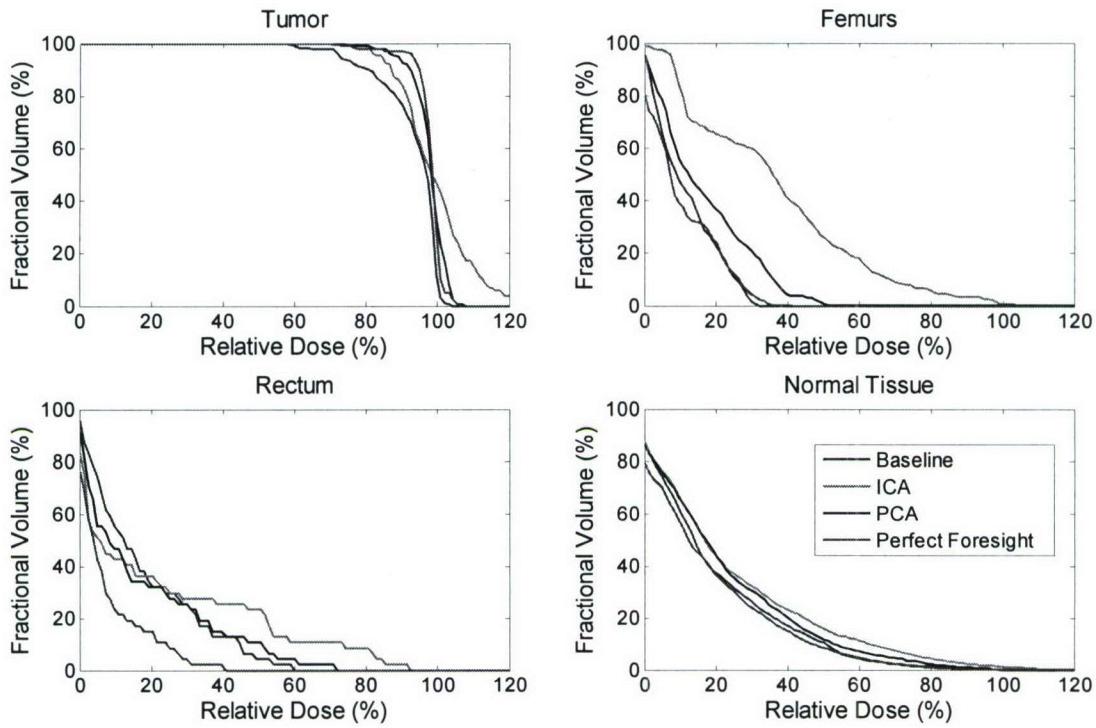


Figure 5b: DVHs of cumulative dose delivered in prostate study as percent of prescription

Algorithm	Tumor			Femurs			Rectum			Normal Tissue		
	Avg. Dose	1 σ	Avg. Dose	1 σ								
ICA	+6%	+14%	+196%	+165%	+21%	+73%	+24%	+28%				
PCA	+5%	-44%	+40%	+45%	-5%	+18%	+13%	+11%				
Perfect Foresight	+5%	-49%	-8%	+12%	-54%	-40%	-7%	+2%				

Table 4: Cumulative dose delivered in prostate study for various algorithms relative to the baseline algorithm, *Adapting to Changing Geometry*.

As in the phantom case, the dose distributions (80% and 95% contour lines of Fig. 5a) show that PCA improves upon the baseline by being able to achieve dose escalation and increased uniformity to the tumor while keeping the dose to the rectum (the sensitive structure) low. In contrast, ICA performs even worse than in the phantom case (it delivers more dose to the femurs and normal tissue). Even more than in the phantom

case, the *Perfect Foresight* DVH shows how much the *Adapting to Geometry and Delivered Dose* algorithms can still improve. In particular, the dose could be lower to the femurs and to a lesser extent to normal tissue and the rectum.

4. Discussion

In current clinical practice, patient setup relies primarily on information from simulation and treatment planning(Balter, 2003). During the whole course of treatment, usually the same treatment plan and setup DRRs (digitally reconstructed radiographs) are employed. Clinically, effort is focused on reproducing, with the aid of orthogonal planar images, the patient's geometry at the simulation stage using translations (and occasionally rotations). While this approach is justifiable for treatment of certain types of diseases such as brain tumors, it generally compromises the treatment because inter-fraction variations in volumes and shapes of the target and sensitive structures are not taken into account. Generally, these inter-fraction changes are multi-dimensional because organs can move relative to each other, and in an extreme situation, each voxel within a soft organ can move relative to other voxels in a complicated manner. Deformable image registration (Court and Dong, 2003; Court et al., 2005; Schreibmann et al., 2005) helps by making these inter-fraction changes visible. But even then, compromised treatment is inevitable because the few degrees of freedom in patient setup (translation and rotation) cannot completely correct for the multi-dimensional changes in the patient geometry: patient setup cannot simultaneously align all the involved structures.

ART solves the problems described in the previous paragraph by adjusting every fraction not only the setup but also the treatment plan. We treat ART as a closed-loop control system where every fraction we re-optimize the radiation beams based on the

latest information (e.g., coming from CBCT). The hope is that the many degrees of freedom available in selecting a plan allow us to compensate for the multi-dimensional changes in the patient geometry. In ART the plan is updated routinely, and the fraction-to-fraction variations of anatomy/physiology and dose delivery lead to modifications of the voxel-specific dose prescription.

ART can be implemented at different levels where the beams can be made to accommodate (i) the new patient setup and the deformed target shape; or (ii) positional/anatomic/physiological changes of all involved organs; or (iii) deformable changes of organs and accumulated dose delivery errors. In conventional 3D conformal radiation therapy, it is not uncommon for a physician to modify a beam portal under the guidance of portal films while the patient is on the treatment couch. In a sense, this is an example of the first kind of ART listed above. However to adapt well to multi-dimensional organ deformations, the number of variables in the beam should be large. Thus the modality of choice for the treatments of type (ii) or (iii) is IMRT or alike. The information available in a particular treatment setup (e.g., whether information is captured about the delivered dose) restricts the types of ART possible in that setup. Our results show that correcting for dose delivery errors (i.e., type iii ART) incurred in previous treatment fractions is important. In conventional radiation therapy (and type i and ii ART) dose errors accumulate because they lack such a compensating mechanism.

ART is a new strategy and its implementation entails the development of a number of “enabling” tools, such as deformable image registration, automated contour mapping from planning CT to the CBCT just before treatment, and effective dose optimization algorithms. This paper tackles the dose optimization problem in ART by developing

suitable dynamic control algorithms. These algorithms try to rapidly optimize the treatment plan each time a new set of input data is available. We propose two types of closed-loop control algorithms (*Adapting to Changing Geometry* and *Adapting to Geometry and Delivered Dose*, see Table 2 for details) for different clinical applications. *Adapting to Changing Geometry* is useful when the accumulated dose is not known accurately. At this point, the deformable registration, which is essential for the calculation of accumulative dose, is still not robust enough for clinical use. *Adapting to Changing Geometry* is particularly helpful in this situation as it affords a currently implementable technique to cope with the nuisance caused by organ deformation. The “manual” approach described by Mohan et al (Mohan et al., 2005) represents a special example of geometric adaptation. The algorithms *Adapting to Geometry and Delivered Dose* are designed to optimize the radiation treatment when both geometric and dosimetric updates are available from time to time. The two algorithms *Adapting to Geometry and Delivered Dose* (ICA and PCA) differ in how they use the update data to direct the treatment. First, ICA is proposed for the situation where the dose compensation needs to be performed right after each system parameters update. Clinically, ICA is valuable in eliminating the unpredictable daily changes resulting from random and fraction specific organ deformation, rectum or bladder filling. Disease sites that are likely to benefit from the treatment include, but are not limited to, prostate, rectum, and cervical cancers. In PCA the task of dose compensation is accomplished by spreading the previous dosimetric errors over a number of subsequent fractions. For practical or clinical reason, the volumetric images just before treatment may not be available on daily basis. In this case, PCA becomes a viable option for certain types of diseases since it distributes the task of compensating a dosimetric error among a number of fractions instead of all-in-one. Indeed, there are clinical situations where the variation of the anatomy may not be notable from day to day but over a larger time span.

Tumor shrinkage and weight loss in some head and neck cases represents a model example, where the CBCT and re-planning do not need to be performed on daily basis. Re-planning using PCA would allow us in this case to maximally benefit from state-of-the-art imaging information. In a way, PCA is similar to the off-line “dumped” correction scheme proposed by Mackie et al (Keller et al., 2003; Mackie et al., 2003), in which the correction is done at a certain point(s) of the treatment course and then applied to several subsequent fractions.

Although showing significant improvement, we note that the algorithms described in this paper are still far from the maximum theoretical performance limit. We encourage the development of more sophisticated closed-loop control algorithms that come closer to this limit. Better predictive schemes for future patient geometry should improve the PCA algorithm greatly. Our use of a weighted quadratic objective function for plan optimization (specifically the selection of beamlet weights) is probably suboptimal (Bedford and Webb, 2006; Cao et al., 2006; Choi and Deasy, 2002; Jeraj et al., 2003; Popple et al., 2005; Thieke et al., 2003; Wu et al., 2002; Yang and Xing, 2004a, , 2004b). However, all the algorithms discussed in this paper would work with other objective functions. One could probably construct an optimal algorithm using the theory of optimal stochastic control. One hurdle is the fact that after delivering the treatment plan for some fraction we must specify the probability of every possible system state. We are researching these ideas.

Finally, we emphasize that the proposed ART is purely dose based and doesn't consider any radiobiological effects. In principle treatment plan optimization should be based on biological models as they are clinically the most relevant. In practice however, there is much controversy to this approach: the dose-response function linking the biological

effect to the radiation dose is not sufficiently understood for various structures. Yang and Xing(Yang and Xing, 2005) have recently proposed a general time-dose-fractionation optimization strategy. This idea combined with our ART framework could potentially handle biologically adaptive radiation therapy. This is clearly a subject of future research.

5. Conclusion

We have proposed a novel ART treatment planning framework to better exploit the information from on-board imaging devices. This dynamic closed-loop control framework is a natural fit for ART and optimizes the dose delivery over the entire course of treatment. We presented three ART algorithms that fit in the closed-loop control framework: one *Adapting to Changing Geometry* and two *Adapting to Geometry and Delivered Dose* (ICA and PCA). All three algorithms incorporate volumetric imaging information just before treatment and the latter two also incorporate information on the accumulated dose. Application of the proposed algorithms to the phantom and clinical cases indicates that the algorithms using information about the accumulated dose (ICA and PCA) allow an escalation of the dose to the tumor. Overall, PCA performed best. Clinical implementation of closed-loop control algorithms should lead to a positive impact in cancer radiotherapy. Given the strong correlation between the doses delivered to the patient and clinical outcome, this new paradigm of ART holds significant promise for us to improve upon the current practice.

Acknowledgement

This work was supported in part by grants from the Department of Defense (PC040282), the National Cancer Institute (1R01 CA104205), the American Cancer Society (RSG-01-022-01-CCE), and a Graduate Research Fellowship from the National Science

Foundation.

References

- Alber M, Nusslin F. 1999. An objective function for radiation treatment optimization based on local biological measures. *Physics in Medicine & Biology* 44:479-493.
- Balter J. 2003. Target and Critical Structure Definitions, Dose Prescription and Re porting for IMRT. In: Palta J, Mackie T, editors. *Intensity-Modulated Radiation Therapy: The State of the Art*. Colorado Springs, CO: Medical Physics Publishing. p 183-198.
- Bedford JL, Webb S. 2003. Elimination of importance factors for clinically accurate selection of beam orientations, beam weights and wedge angles in conformal radiation therapy. *Med Phys* 30:1788-1804.
- Bedford JL, Webb S. 2006. Constrained segment shapes in direct-aperture optimization for step-and-shoot IMRT. *Med Phys* 33:944-958.
- Bortfeld T. 1999. Optimized planning using physical objectives and constraints. *Seminars in Radiation Oncology* 9:20-34.
- Cao D, Earl MA, Luan S, Shepard DM. 2006. Continuous intensity map optimization (CIMO): a novel approach to leaf sequencing in step and shoot IMRT. *Med Phys* 33:859-867.
- Censor Y. 2003. Mathematical Optimization for the Inverse Problem of Intensity-Modulated Radiation Therapy. In: Palta J, Mackie T, editors. *Intensity-Modulated Radiation Therapy: The State of the Art*. Colorado Springs, CO: Medical Physics Publishing. p 25-50.
- Choi B, Deasy JO. 2002. The generalized equivalent uniform dose function as a basis for intensity-modulated treatment planning. *Phys Med Biol* 47:3579-3589.
- Corletto D, Iori M, Paiusco M, Brait L, Broggi S, Ceresoli G, Iotti C, Calandrino R, Fiorino C. 2003. Inverse and forward optimization of one- and two-dimensional intensity-modulated radiation therapy-based treatment of concave-shaped planning target volumes: the case of prostate cancer. *Radiother Oncol* 66:185-195.
- Court LE, Dong L. 2003. Automatic registration of the prostate for computed-tomography-guided radiotherapy. *Medical Physics* 30:2750-2757.
- Court LE, Dong L, Lee AK, Cheung R, Bonnen MD, O'Daniel J, Wang H, Mohan R, Kuban D. 2005. An automatic CT-guided adaptive radiation therapy technique by online modification of multileaf collimator leaf positions for prostate cancer. *Int J Radiat Oncol Biol Phys* 62:154-163.
- de la Zerda A, Ambrush B, Xing L. 2006. Inverse planning for adaptive radiation therapy using dynamic programming. In: Annual Meeting of ASTRO. Philadelphia, PA. p 2163.
- Dietrich L, Jetter S, Tucking T, Nill S, Oelfke U. 2006. Linac-integrated 4D cone beam CT: first experimental results. *Phys Med Biol* 51:2939-2952.
- Gill PE, Murray W, Saunders MA. 2005. SNOP: an SQP algorithm for large-scale constrained optimization. *SIAM Review* 47:99-131.
- Hristov DH, Fallone BG. 1997. An active set algorithm for treatment planning optimization. *Medical Physics* 24:1455-1464.
- Jeraj R, Wu C, Mackie TR. 2003. Optimizer convergence and local minima errors and their clinical importance. *Phys Med Biol* 48:2809-2827.
- Keller H, Ritter MA, Mackie TR. 2003. Optimal stochastic correction strategies for rigid-body target motion. *Int J Radiat Oncol Biol Phys* 55:261-270.
- Langen KM, Meeks SL, Poole DO, Wagner TH, Willoughby TR, Kupelian PA, Ruchala KJ, Haimerl J, Olivera GH. 2005. The use of megavoltage CT (MVCT) images for dose recomputations. *Phys Med Biol* 50:4259-4276.

- Li RP, Yin FF. 2000. Optimization of inverse treatment planning using a fuzzy weight function. *Medical Physics* 27:691-700.
- Li T, Xing L, McGuinness C, Munro P, Loo B, Koong A. 2006. Four-dimensional cone-beam CT using an on-board imager. *Medical Physics* 33:3825-3833.
- Mackie TR, Kapatoes J, Ruchala K, Lu W, Wu C, Olivera G, Forrest L, Tome W, Welsh J, Jeraj R, Harari P, Reckwerdt P, Paliwal B, Ritter M, Keller H, Fowler J, Mehta M. 2003. Image guidance for precise conformal radiotherapy. *Int J Radiat Oncol Biol Phys* 56:89-105.
- Mohan R, Zhang X, Wang H, Kang Y, Wang X, Liu H, Ang KK, Kuban D, Dong L. 2005. Use of deformed intensity distributions for on-line modification of image-guided IMRT to account for interfractional anatomic changes. *Int J Radiat Oncol Biol Phys* 61:1258-1266.
- Oelfke U, Bortfeld T. 1999. Inverse planning for x-ray rotation therapy: a general solution of the inverse problem. *Physics in Medicine & Biology* 44:1089-1104.
- Oldham M, Letourneau D, Watt L, Hugo G, Yan D, Lockman D, Kim LH, Chen PY, Martinez A, Wong JW. 2005. Cone-beam-CT guided radiation therapy: A model for on-line application. *Radiother Oncol* 75:271 E271-278.
- Olivera GH, Mackie TR, Ruchala K, Lu W, Kapatoes J. 2006. Adaptive radiation therapy (art) strategies using helical tomotherapy. In: Bortfeld T, Schmidt-Ullrich R, De Deve W, Wazer DE, editors. *Image-Guided IMRT*. Berlin: Springer-Verlag. p 235-246.
- Popple RA, Prelop PB, Spencer SA, De Los Santos JF, Duan J, Fiveash JB, Brezovich IA. 2005. Simultaneous optimization of sequential IMRT plans. *Med Phys* 32:3257-3266.
- Pouliot J, Bani-Hashemi A, Chen J, Svatos M, Ghelmansarai F, Mitschke M, Aubin M, Xia P, Morin O, Bucci K, Roach M, 3rd, Hernandez P, Zheng Z, Hristov D, Verhey L. 2005. Low-dose megavoltage cone-beam CT for radiation therapy. *Int J Radiat Oncol Biol Phys* 61:552-560.
- Rowbottom CG, Nutting CM, Webb S. 2001. Beam-orientation optimization of intensity-modulated radiotherapy: clinical application to parotid gland tumours. *Radiotherapy & Oncology* 59:169-177.
- Schreibmann E, Yang Y, Boyer A, Li T, Xing L. 2005. Image Interpolation in 4D CT Using a BSpline Deformable Registration Model. *Medical Physics* 32:1924.
- Shepard DM, Earl MA, Li XA, Naqvi S, Yu C. 2002. Direct aperture optimization: a turnkey solution for step-and-shoot IMRT. *Medical Physics* 29:1007-1018.
- Soderstrom S, Gustafsson A, Brahme A. 1993. The clinical value of different treatment objectives and degrees of freedom in radiation therapy optimization. *Radiotherapy & Oncology* 29:148-163.
- Sonke JJ, Zijp L, Remeijer P, van Herk M. 2005. Respiratory correlated cone beam CT. *Med Phys* 32:1176-1186.
- Thieke C, Bortfeld T, Niemierko A, Nill S. 2003. From physical dose constraints to equivalent uniform dose constraints in inverse radiotherapy planning. *Med Phys* 30:2332-2339.
- Trofimov A, Reitzel E, Lu HM, Martin B, Jiang S, Chen G, Bortfeld T. 2005. Temporo-spatial IMRT optimization: concepts, implementation and initial results. *Physics in Medicine and Biology* 50:in press.
- Webb S, Convery DJ, Evans PM. 1998. Inverse planning with constraints to generate smoothed intensity-modulated beams. *Physics in Medicine & Biology* 43:2785-2794.
- Widrow B, Stearns S. 1985. *Adaptive Signal Processing*. Upper Saddle River, NJ: Prentice Hall.
- Widrow B, Walach E. 1995. *Adaptive Inverse Control*. Upper Saddle River, NJ: Prentice Hall.
- Wu Q, Liang J, Yan D. 2006. Application of dose compensation in image-guided radiotherapy of prostate cancer. *Phys Med Biol* 51:1405-1419.
- Wu Q, Mohan R, Niemierko A, Schmidt-Ullrich R. 2002. Optimization of intensity-modulated radiotherapy plans based on the equivalent uniform dose. *International Journal of Radiation Oncology, Biology, Physics* 52:224-235.

- Xiao Y, Galvin J, Hossain M, Valicenti R. 2000. An optimized forward-planning technique for intensity modulated radiation therapy. *Medical Physics* 27:2093-2099.
- Xing L, Li JG, Donaldson S, Le QT, Boyer AL. 1999. Optimization of importance factors in inverse planning. *Physics in Medicine & Biology* 44:2525-2536.
- Xing L, Wu Q, Yang Y, Boyer A. 2005. Physics of IMRT. In: Mundt A, Roeske J, editors. *Intensity-Modulated Radiation Therapy: A Clinical Perspective*. Hamilton, London: BC Decker Inc. p 20-52.
- Yan D, Wong J, Vicini F, Michalski J, Pan C, Frazier A, Horwitz E, Martinez A. 1997. Adaptive modification of treatment planning to minimize the deleterious effects of treatment setup errors. *Int J Radiat Oncol Biol Phys* 38:197-206.
- Yang Y, Schreibmann E, Li T, Xing L. 2006. Dosimetric evaluation of kV cone-beam CT (CBCT) based dose calculation. *Medical Physics* submitted.
- Yang Y, Xing L. 2004a. Clinical knowledge-based inverse treatment planning. *Phys Med Biol* 49:5101-5117.
- Yang Y, Xing L. 2004b. Inverse treatment planning with adaptively evolving voxel-dependent penalty scheme. *Med Phys* 31:2839-2844.
- Yang Y, Xing L. 2005. Optimization of radiation dose-time-fractionation scheme with consideration of tumor specific biology. *Medical Physics* 32:3666-3677.

MULTISCALE IMAGE REGISTRATION

DANA PAQUIN

Department of Mathematics, Stanford University
Stanford, CA 94305-2125

DORON LEVY

Department of Mathematics, Stanford University
Stanford, CA 94305-2125

EDUARD SCHREIBMANN

Department of Radiation Oncology, Stanford University
Stanford, CA 94305-5947

LEI XING

Department of Radiation Oncology, Stanford University
Stanford, CA 94305-5947

(Communicated by Yang Kuang)

ABSTRACT. A multiscale image registration technique is presented for the registration of medical images that contain significant levels of noise. An overview of the medical image registration problem is presented, and various registration techniques are discussed. Experiments using mean squares, normalized correlation, and mutual information optimal linear registration are presented that determine the noise levels at which registration using these techniques fails. Further experiments in which classical denoising algorithms are applied prior to registration are presented, and it is shown that registration fails in this case for significantly high levels of noise, as well. The hierarchical multiscale image decomposition of E. Tadmor, S. Nezzar, and L. Vese [20] is presented, and accurate registration of noisy images is achieved by obtaining a hierarchical multiscale decomposition of the images and registering the resulting components. This approach enables successful registration of images that contain noise levels well beyond the level at which ordinary optimal linear registration fails. Image registration experiments demonstrate the accuracy and efficiency of the multiscale registration technique, and for all noise levels, the multiscale technique is as accurate as or more accurate than ordinary registration techniques.

1. Introduction. Often in image processing, images must be spatially aligned to allow practitioners to perform quantitative analyses of the images. The process of aligning images taken, for example, at different times, from different imaging devices, or from different perspectives, is called image registration. More precisely, image registration is the process of determining the optimal spatial transformation that maps one image to another. Typically, two images are taken as input, and the registration process is then the optimization problem which determines the

2000 *Mathematics Subject Classification.* Primary: 68U10; Secondary: 92C55, 62P10, 94A08.

Key words and phrases. image registration, multiscale analysis, noise, CT, MRI, mutual information.

geometric mapping that brings one image into spatial alignment with the other image. In practice, the particular type of transformation as well as the notion of optimal will depend on the specific application.

Examples of applications in which image registration is particularly important include astro- and geophysics, computer vision, remote sensing, and medicine. In this paper, we will focus on medical image registration. Image registration plays an important role in the analysis of medical images. For example, images taken from different sensors often contain complementary information. By bringing the two images into alignment so that anatomical features of one modality can be detected in the other modality, the information from the different modalities can be combined. In neurosurgery, for example, tumors are typically identified and diagnosed using magnetic resonance images (MRI), but stereotaxy technology (the use of surgical instruments to reach specified points) generally uses computed tomography (CT) images. Registration of these modalities allows the transfer of coordinates of tumors from the MRI images to the CT images. See [14] for a discussion of the applications of multimodality imaging to problems in neurosurgery. As another example, medical image data acquired prior to diagnosis can be compared with data acquired during or after treatment to determine the effectiveness of the treatment. To compare images taken at different times, however, the images must first be brought into spatial alignment so that actual differences in the data can be distinguished from differences that result from the image acquisition process.

In the context of medical imaging, the goal of the registration process is to remove artificial differences in the images introduced by patient movement, differences in imaging devices, etc., but at the same time, to retain real differences due to actual variations of the objects. Medical images, however, often contain significant levels of noise due to instrumentation imperfections, data acquisition techniques, image reconstruction methods, transmission and/or compression errors, and other factors. Although numerous successful image registration techniques have been published, we will see that ordinary image registration algorithms can fail to produce meaningful results when one or both of the images to be registered contains significant levels of noise.

Since noise is generally present in digital images, image denoising is a fundamental problem in image processing. Indeed, many approaches to image denoising have been presented. Thus a simple solution to the problem of image registration in the presence of noise would be to first apply a denoising algorithm to the noisy image(s), and then use existing image registration techniques to register the denoised images. However, common denoising algorithms, most notably spatial filtering algorithms, have the disadvantage that while they are successful in removing noise, they often remove edges as well. Additionally, most denoising procedures require a priori knowledge of the noise level, variance, and/or model, information not typically known in practice. For these and other reasons, we will demonstrate that ordinary image registration of noisy images fails to produce acceptable results even when classical denoising algorithms are applied to the noisy images prior to registration (for significantly high levels of noise). Moreover, even more advanced denoising techniques such as anisotropic diffusion (which was designed to remove noise while preserving edges) will be shown to fail to register noisy images. Thus, we seek a technique that enables successful image registration when one or both of the images to be registered is noisy.

Generally, we would like to consider an image f consisting of *coarse* and *fine* scales. The general shape and main features of an image are considered the coarse scales, and details and textures, such as noise, are the fine scales of the image. Separating the coarse and fine scales of an image, therefore, is an effective tool in denoising. Indeed, several denoising algorithms have been proposed using separation of the coarse and fine scales of an image, most notably [19], [18], [11], and [20]. The method presented in [20] presents a multiscale technique in which an image f is decomposed in a hierarchical expansion $f \sim \Sigma_j u_j$, where the u_j (called the components of f relative to the decomposition) resolve edges of f with increasing scales. More precisely, for small k , the sum $\Sigma_j^k u_j$ is a coarse representation of the image f , and as k increases, the sum captures more and more detail (and hence, noise) of the image.

In this paper, we present a multiscale image registration technique based on the multiscale decomposition of [20] that is particularly effective when one or both of the images to be registered contains significant levels of noise. Since the hierarchical expansion $f \sim \Sigma_j u_j$ decomposes the image f into components which contain increasingly fine scales, we expect a component-wise registration algorithm to produce accurate results for noisy images. That is, given a noisy image f , for small values of k , the component $\Sigma_j^k u_j$ retains the general shape of the image f while removing the details and noise of the image. Thus, if we wish to register the image f with another image, say g , we expect that registration of the components $\Sigma_j^k u_j$ with g will provide an accurate estimation of the actual transformation that brings the two images into spatial alignment with one another, for sufficiently small values of k . Similarly, if both f and g are noisy, we expect decomposing both images and performing component-wise registrations should accurately estimate the optimal transformation. We will demonstrate that multiscale image registration enables successful image registration for images that contain levels of noise that are significantly higher than the levels at which ordinary registration fails.

This paper is organized in the following way. In Section 2, we discuss the image registration problem and review standard image registration techniques. In Section 3, we present the problem of image registration in the presence of noise, and illustrate the failure of current techniques when one or both of the images to be registered contains high levels of noise. In Section 4, we briefly discuss classical and modern denoising techniques, and illustrate the failure of ordinary image registration of noisy images even when the images are denoised prior to registration. In Section 5, we review the multiscale image decomposition of [20], and illustrate the results of the hierarchical multiscale decomposition obtained upon applying the algorithm to noisy images. In Section 6, we present image registration techniques based upon the multiscale decomposition described in Section 5, and in Section 7, we present the results of our multiscale image registration experiments.

2. The registration problem. Given a *fixed* and *moving* image, the registration problem is the process of finding an *optimal transformation* that brings the moving image into spatial alignment with the fixed image. While this problem is easy to state, it is difficult to solve. The main source of difficulty is that the problem is ill-posed, which means, for example, that the problem may not have a unique solution. Additionally, the notion of optimality may vary for each application: for example, some applications may require consideration only of rigid transformations, while other applications require non-rigid transformations. Finally, computation time

and data storage constraints place limitations on the complexity of models that can be used for describing the problem. This following discussion of image registration follows the presentation in [12].

2.1. The mathematical setting. A two-dimensional gray-scale image f is a mapping which assigns to every point $x \in \Omega \subset \mathbb{R}^2$ a gray value $f(x)$ (called the intensity value of the image at the point x). We will consider images as elements of the space $L^2(\mathbb{R}^2)$. Color images can be defined, for example, in terms of vector-valued functions $\mathbf{f} = (f_1, f_2, f_3)$ representing the RGB-color scales. For the medical imaging applications that we are interested in, images are in fact given in terms of discrete data, and the function f must be obtained by interpolation. We will not discuss this construction here, but we will assume that an interpolation method has been chosen.

Image registration is necessary, for example, for two images of the same object that are not spatially aligned. This occurs when the images are taken at different times, from different perspectives, or from different imaging devices. The basic input data to the registration process are two images: one is defined as the fixed image $f(x)$ and the other as the moving image $m(x)$. The goal is then to find a transformation ϕ such that the fixed image $f(x)$ and the deformed moving image $m_\phi(x) := m(\phi(x))$ are similar. To solve this problem in a mathematical way, the term *similar* needs to be defined appropriately. For example, if the images to be registered are taken from different devices, there may not be a correspondence between the intensities $f(x)$ and $m_\phi(x)$ for an optimal ϕ . Additionally, we may consider measures of similarity between the images which are not related to the intensities. Thus, the registration problem necessarily involves a discussion of the distance measures, or metrics, used to compare images. The general problem can then be stated as follows.

Given a distance measure $D : (L^2(\mathbb{R}^2))^2 \rightarrow \mathbb{R}$ and two images $f(x), m(x) \in L^2(\mathbb{R}^2)$, the solution ϕ of the registration problem is given by the following minimization problem:

$$\phi = \operatorname{argmin}_{\psi: \mathbb{R}^2 \rightarrow \mathbb{R}^2} D(f, m_\psi). \quad (1)$$

In many applications, the set of allowable transformations to be considered in the minimization problem (1) is restricted to a strict subset of the set of all maps $\psi : \mathbb{R}^2 \rightarrow \mathbb{R}^2$. For example, we may require the transformation ϕ to be smooth, or we may impose specific parametric requirements, such as requiring the transformation to be rigid, affine, polynomial, etc.

2.2. Landmark-based registration. Landmark-based registration is an image registration technique based on a finite set of image features. The problem is to determine the transformation such that for a finite set of features, any feature of the moving image is mapped onto the corresponding features of the fixed image. More precisely, let $F(f, j)$ and $F(m, j)$, $j = 1, \dots, m$ be given features of the fixed and moving images, respectively. The solution ϕ of the registration problem is then a map $\phi : \mathbb{R}^2 \rightarrow \mathbb{R}^2$ such that

$$F(f, j) = \phi(F(m, j)), \quad j = 1, \dots, m. \quad (2)$$

For a more general notion of landmark-based registration, we define the following distance measure:

$$D^{LM}(\phi) := \sum_{j=1}^m \|F(f, j) - \phi(F(m, j))\|_l^2, \quad (3)$$

where $\|\cdot\|_l$ denotes a norm on the landmark, or feature, space. For example, if the features are locations of points, then $\|\cdot\|_l = \|\cdot\|_{\mathbb{R}^2}$. We can then restate (2) as the minimization problem in which the solution $\phi : \mathbb{R}^2 \rightarrow \mathbb{R}^2$ of the registration problem is given by:

$$\phi = \operatorname{argmin}_{\psi: \mathbb{R}^2 \rightarrow \mathbb{R}^2} D^{LM}(\psi). \quad (4)$$

To solve this minimization problem, the transformation either is chosen to be an element of an n -dimensional space spanned, for example, by polynomials, splines, or wavelets, or it is required to be smooth in some sense. In the first case, the features to be mapped are the locations of a number of user-supplied landmarks. Let χ_k , $k = 1, \dots, n$ be the basis functions of the space. Then the minimization of

$$D^{LM}(\phi) := \sum_{j=1}^m \|F(f, j) - \phi(F(m, j))\|_l^2$$

can be obtained upon expanding $\phi = (\phi_1, \phi_2)$ in terms of the basis functions χ_k and solving the resulting least squares problems.

In the case in which we require the transformation ϕ to be smooth, we introduce a functional which imposes smoothness restrictions on the transformation. That is, we look for a transformation ϕ which interpolates the features $F(f, j)$ and $F(m, j)$, and which is smooth in some sense. Such a transformation is called a minimal norm solution, and it turns out (see [8]) that the solution can be expressed in terms of radial basis functions.

Landmark-based registration is simple to implement, and the numerical solution requires only the solution of a linear system of equations. However, the main drawback of the landmark-based approach is that the registration process depends on the location of the landmarks. As the detection and mathematical characterization of landmarks (for example, anatomical landmarks in medical images) is not fully automated, the landmarks must be user-supplied, and this can be a time-consuming and difficult process, even for a medical expert; see, for example, [17]. Additionally, landmark-based registration does not always result in a physically meaningful registration. See [12, p. 44], for a simple example of a situation in which landmark-based registration fails to produce meaningful results.

2.3. Principal-axes-based registration. Principal-axes image registration is based on the idea of landmark-based registration, but it uses features that can be automatically detected. These features are constructed as follows. For an image $f : \mathbb{R}^2 \rightarrow \mathbb{R}$, and a function $g : \mathbb{R}^2 \rightarrow \mathbb{R}$, we define the expectation value of g with respect to f by

$$\mathbb{E}_f(g) := \frac{\int_{\mathbb{R}^2} g(x)f(x) dx}{\int_{\mathbb{R}^2} f(x) dx}. \quad (5)$$

If $u : \mathbb{R}^2 \rightarrow \mathbb{R}^{m \times n}$, we set $\mathbb{E}_f(u) := \mathbb{E}_f[u_{j,k}] \in \mathbb{R}^{m \times n}$. The center of an image f is defined by

$$C_f := \mathbb{E}_f[x] \in \mathbb{R}^2, \quad (6)$$

and the covariance by

$$\text{Cov}_f := \mathbb{E}_f[(x - C_f)(x - C_f)^T] \in \mathbb{R}^{2 \times 2}. \quad (7)$$

Given fixed and moving images, $f(x)$ and $m(x)$, the centers c_f and c_m and eigendecompositions of the covariance matrices Cov_f and Cov_m are used as the features F_i , and the registration problem is to compute $\phi : \mathbb{R}^2 \rightarrow \mathbb{R}^2$ such that $F_i(m(\phi)) = F_i(f)$ for the features F_i .

This method is described in detail in [1]. The principal-axes method of image registration has the advantages that it is computationally fast and simple and requires few registration parameters, but has the disadvantages that it is not suitable for images of multiple modalities and that the solutions may be ambiguous. In particular, the principal-axes-based method cannot distinguish between images with the same center and covariance, even though images with very different structure and orientation may have the same center and/or covariance.

2.4. Optimal parametric registration. An alternative approach to registration is to use methods that are based on the minimization (or maximization) of some distance measure, or metric, D . The transformation ϕ is restricted to some parameterized space, and the registration can be obtained by minimizing (or maximizing) the distance D over the parameterized space. In particular, we will discuss metrics based on intensity, correlation, and mutual information. Given a metric D , a fixed image f , and a moving image m , optimal parametric registration is the problem of finding a transformation ϕ in some pre-specified parameterizable space such that $D(f, m(\phi))$ is minimized (or maximized in certain cases). Examples of commonly used parameterizable spaces in image registration are polynomial and spline spaces. We will primarily be interested in rigid and affine linear transformations. An affine linear map is a map of the form $\phi(x) = Ax + b$, $A \in \mathbb{R}^{2 \times 2}$, $\det A > 0$, $b \in \mathbb{R}^2$. Such a map allows rotations, translations, scales, and shears of the coordinates. A translation (or rigid) transformation is a special case of an affine transformation which allows only rotations and translations of the coordinates, and in this case, the matrix A is required to be orthogonal with determinant 1. Optimal parametric registration is probably the most commonly used image registration technique.

To minimize $D(f, m(\phi))$, we must choose an optimization technique. That is, an optimal parametric registration technique is described by a metric to be minimized (or maximized) and an optimizer which controls the minimization (or maximization). The implementation of the registration algorithm works in the following way: at each iteration, the distance D between the two images is computed. An affine transformation is then applied to the moving image, and the distance between the images is recomputed. In theory, this process continues until the distance is minimized (or maximized), though in practice there is some stopping criterion.

At each stage, the optimizer determines the parameters of the transformation that will be applied to the moving image. Examples of commonly used optimizers include gradient descent and regular step gradient descent. Gradient descent optimization advances the parameters of the transformation in the direction of the gradient, where the step size is governed by a user-specified learning rate. Regular step gradient descent optimization advances the parameters of the transformation in the direction of the gradient where a bipartition scheme is used to compute the step size.

2.4.1. *The mean squares metric.* The mean squares metric computes the mean-squared pixel-wise difference in intensity between two images f and m :

$$MS(f, m) := \frac{1}{N} \sum_{i=1}^N (f_i - m_i)^2, \quad (8)$$

where N is the total number of pixels considered, f_i is the i^{th} pixel of image f , and m_i is the i^{th} pixel of image m . Note that the optimum value of the mean squares metric is 0, and poor matches between the images f and m result in large values of $MS(f, m)$. This metric has the advantage that it is computationally simple. It is based on the assumption that pixels in one image should have the same intensity as (spatially) corresponding pixels in the second image. Thus, the mean squares metric is restricted in practice to images of the same modality.

2.4.2. *The normalized correlation metric.* The normalized correlation metric computes pixel-wise cross-correlation and normalizes it by the square root of the auto-correlation function:

$$NC(f, m) := -\frac{\sum_{i=1}^N (f_i \cdot m_i)}{\sqrt{\sum_{i=1}^N f_i^2 \cdot \sum_{i=1}^N m_i^2}}, \quad (9)$$

where N , f_i , and m_i are as defined for the mean squares metric. The negative sign in (9) causes the optimum value of the metric to occur when the minimum is reached. Thus the optimal value of the normalized correlation metric is -1. As with the mean squares metric, the normalized correlation metric is restricted to images of the same modality.

2.4.3. *The mutual information metric.* Mutual information is an information-theoretic approach to image registration that was proposed independently by Viola and Wells [22] and Collignon et al. [4] in 1995. The idea is that mutual information computes the amount of information that one random variable (here, image intensity) gives about another random variable (here, intensity values of another image). More precisely, given a fixed image $f(x)$ and a moving image $m(x)$, we wish to compute the transformation ϕ which *maximizes* the mutual information; i.e.,

$$\phi = \arg \max_{\psi} I(f(x), m(\phi(x))). \quad (10)$$

Maximization of the mutual information criterion assumes that the statistical dependence between corresponding image intensity values is maximized when the images are geometrically aligned.

The mutual information $I(f(x), m(\phi(x)))$ is defined in terms of entropy, where we consider x as a random variable over coordinate locations in the coordinate system of the fixed image. Let $h(\cdot)$ denote the entropy of a random variable: $h(x) := - \int p(x) \ln p(x) dx$, where $p(x)$ is the probability density function of the random variable x . Note that it is not clear how to construct $p(x)$; we will discuss methods for estimating the probability densities. The joint entropy of two random variables x and y is given by $h(x, y) = - \int p(x, y) \ln p(x, y) dx dy$, where $p(x, y)$ is the joint probability density function of the random variables x and y . Entropy can be considered as a measure of the uncertainty or complexity of a random variable.

If x and y are independent, then $p(x, y) = p(x)p(y)$, so $h(x, y) = h(x) + h(y)$. However, if there is any dependency (as would be the case if x and y are intensity values of images of the same object), then $h(x, y) < h(x) + h(y)$. The difference is defined to be *mutual information*:

$$I(f(x), m(\phi(x))) = h(f(x)) + h(m(\phi(x))) - h(f(x), m(\phi(x))). \quad (11)$$

The terms in (11) can be interpreted in the following way. The first term, $h(f(x))$, is the entropy of the fixed image and is independent of the transformation ϕ . The second term, $h(m(\phi(x)))$, is the entropy of $m(\phi(x))$, so maximization of mutual information encourages transformations ϕ for which $m(\phi(x))$ has a high level of complexity or uncertainty. The third term $-h(f(x), m(\phi(x)))$ is the negative joint entropy of $f(x)$ and $m(\phi(x))$, so maximization of mutual information is related to minimization of the joint entropy of $f(x)$ and $m(\phi(x))$. A detailed overview of mutual information based registration can be found in [16].

Mutual information has the following properties. Let $u(x)$ and $v(x)$ denote two images.

1. $I(u(x), v(x)) = I(v(x), u(x))$. Mutual information is symmetric.
2. $I(u(x), u(x)) = h(u(x))$. The information an image contains about itself is equal to the entropy of the image.
3. $I(u(x), v(x)) \leq h(u(x))$ and $I(u(x), v(x)) \leq h(v(x))$. The information that the images contain about each other can not be greater than the information contained in the individual images.
4. $I(u(x), v(x)) \geq 0$.
5. $I(u(x), v(x)) = 0$ if and only if $u(x)$ and $v(x)$ are independent. If the images $u(x)$ and $v(x)$ are independent, no information about one image is gained when the other image is known.

The entropies in equation (11) are defined in terms of integrals over the probability densities associated with the images $f(x)$ and $m(x)$. However, in a typical medical image registration problem, the probability densities are not directly accessible, and thus must be estimated from the image data. Parzen windowing, described in [5] and used in [22], is a common technique for density estimation. In this method, continuous density functions are constructed by a super-position of kernel functions $K(\cdot)$ centered at the elements of a sample of intensities taken from the image. The estimation of the probability density $p(z)$ is thus given by

$$p(x) \cong P^*(z) = \frac{1}{N_S} \sum_{z_j \in S} K(z - z_j), \quad (12)$$

where N_S is the number of spatial samples in S and K is an appropriately chosen kernel function. The kernel function K must be smooth, symmetric, have zero mean, and unit mass. Examples of suitable candidates for K include the Gaussian density and the Cauchy density. In [22], Viola and Wells use a Gaussian density function with standard deviation σ to estimate the probability density functions. The optimal value of σ depends on the particular images to be registered.

Upon estimating the probability densities using the Parzen windowing technique, the entropy integral $h(z) = -\int p(z) \ln(p(z)) dz$ must be evaluated, for example, by using a sample mean:

$$h(z) \cong -\frac{1}{N_R} \sum_{z_j \in R} \ln(P^*(z_j)), \quad (13)$$

where R is a second sample of intensities taken from the image. That is, two separate intensity samples S and R are taken from the image. The first is used to estimate the probability density, and the second is used to approximate the entropy.

The main advantage of the mutual information measure is that was shown to be generally applicable for multi-modality registration, whereas intensity-based measures are typically not applicable for multimodality registration. Mutual information registration has been successfully used for a number of complex applications. Most notably, mutual information has been shown to be highly accurate for MRI-CT registration; see, for example, [9], [15], and [21].

2.5. Non-parametric image registration. All of the image registration techniques that we have discussed so far have been based on certain parameters. For example, either the transformation ϕ can be expanded in terms of basis functions that span a specified finite-dimensional space, or the registration is controlled by a specified set of external features. Non-parametric techniques do not restrict the transformation to a parameterizable set. Given two images, a fixed image $f(x)$ and a moving image $m(x)$, we seek a transformation ϕ such that $m(\phi(x))$ is similar to $f(x)$ in a certain sense. Upon defining a suitable distance measure D , the registration problem is then to minimize the distance between $m(\phi(x))$ and $f(x)$. However, a direct minimization is often not possible in the non-parametric case. The problem is ill-posed: small changes in the input data may lead to large changes in the output. Additionally, the solution is not unique. Given these constraints, a stable numerical implementation is often impossible. To circumvent these problems, a regularizing, or smoothing, term S is introduced, and the registration problem becomes the minimization of the distance between $m(\phi(x))$ and $f(x)$ plus a smoothing term $S(\phi)$. That is, the registration is based on a regularized minimization of the distance between the images.

In the discussion of non-parametric image registration, the transformation $\phi : \mathbb{R}^2 \rightarrow \mathbb{R}^2$ is split into the trivial identity part and the deformation or displacement part u ; i.e.,

$$\phi(x) = x - u(x). \quad (14)$$

Upon decomposing ϕ in this way, we have $m(\phi(x)) = m(x - u(x)) := m_u(x)$. Given a distance D and a smoother S , the elastic registration problem is then the minimization of $D(f(x), m_u(x)) + \alpha S(u)$, where $\alpha \in \mathbb{R}$ is a positive regularizing parameter.

The choice of smoother S typically depends on the particular application. Examples of non-parametric image registration techniques include elastic registration [3], fluid registration [2], and diffusion registration [6]. Elastic registration uses linear elasticity theory to model the deformation of an elastic body. In this case, the regularizing term $S(u)$ is the linearized elastic potential of the displacement u . In fluid registration, the regularization is based on the linearized elastic potential of the time derivative of u . Finally, diffusion registration uses a regularization that is based on spatial derivatives of the displacement.

REMARK. *In this section, we presented a brief overview of the major image registration techniques currently used in image registration. In practice, the best*

registration method for a given set of images will depend on the particular features of the images themselves. However, numerous studies comparing the accuracy and performance of different image registration techniques for various applications have been presented. The most extensive of these is [24], which originally consisted of a comparison of 16 methods but has since been substantially expanded.

3. Registration in the presence of noise. In this section, we study the effect of noise on image registration, and we determine the approximate noise level at which registration fails. This study is conducted on the brain proton density slice images shown in Figure 1 below. The image on the right is the result of translating the image on the left by 13 mm to the right in the X -direction and 17 mm downward in the Y -direction. Let I denote the original image, and let T denote the translated image.

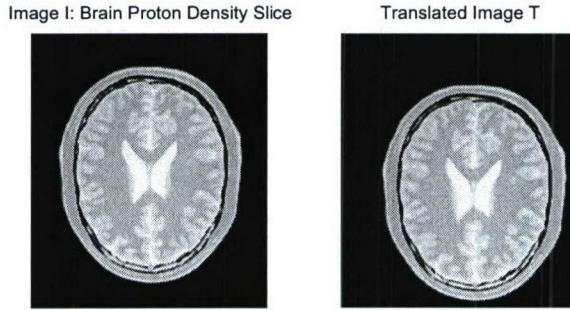


FIGURE 1. Original image I and translated image T .

Initially, we will consider the registration problem in which one of the images (here, the fixed image) is noisy. We will add increasing levels of noise to the image I and register the non-noisy translated image T with the noisy images. Our goal is to determine the approximate noise levels at which various image registration techniques fail, and to develop an algorithm that will enable registration beyond these levels. Since we know the exact transformation that brings T into spatial alignment with I , we can effectively evaluate and compare the accuracy of various registration techniques. We will demonstrate that our multiscale registration technique enables accurate registration of the translated image T with images that contain significant levels of noise. Eventually, we will also apply our techniques to the case in which both the fixed and the moving images contain high levels of noise. Before we present these results, we discuss the notion of noise in some detail.

REMARK. *In this paper we present the results only for registration experiments using the images I and T shown in Figure 1. We have performed numerous experiments using other images, and we obtained results similar to those presented in this paper. For the sake of brevity, we limit the results presented in this paper to the experiments using the images in Figure 1.*

3.1. Noise. Digital images are often degraded by random noise. In imaging, the term noise refers to random fluctuations in intensity values that occur during image capture, transmission, or processing, and that may distort the information given by

the image. Image noise is not part of the ideal signal and may be caused by a wide range of sources, such as detector sensitivity, environmental radiation, transmission errors, discretization effects, etc. Noise is generally classified as either independent noise or noise which is dependent on the image data.

Independent noise can often be described by an additive noise model, in which the observed image $f(x)$ is the sum of the true image $s(x)$ and the noise $n(x)$:

$$f(x) = s(x) + n(x). \quad (15)$$

Within this framework of additive noise, the noise $n(x)$ is commonly modeled by Gaussian noise of mean m and variance v . A multiplicative noise model describes noise that is dependent on the image data. This is often referred to as speckle noise:

$$f(x) = s(x) + s(x)n(x) = s(x)(1 + n(x)). \quad (16)$$

In this case, $n(x)$ is uniformly distributed random noise with mean m and variance v . Impulse noise, or salt-and-pepper noise, is noise that resembles salt and pepper granules randomly distributed over the image. Impulse noise is typically defined by the following model. We let $s(x)$ denote the actual image, and $f(x)$ denote the observed image. Then

$$f(x) = \begin{cases} s(x), & \text{with probability } 1 - \delta, \\ \eta(x), & \text{with probability } \delta, \end{cases} \quad (17)$$

where $\eta(x)$ is an identically distributed, independent random process. With this model, an arbitrary pixel $x \in \Omega \subset \mathbb{R}^2$ is affected by noise with probability δ , and not affected with probability $1 - \delta$. We will refer to δ as the impulse noise density, as adding impulse noise of density δ to an image $f(x)$ affects approximately $\delta \cdot \text{size}(f)$ pixels. The random process $\eta(x)$ is typically such that the corrupted pixels are either set to the maximum value, have single bits flipped over, or are set alternatively to zero or to the maximum value. This last case results in a salt-and-pepper appearance. Note that unaffected pixels always remain unchanged.

In Figure 2, we add additive Gaussian noise of mean 0 and variance 0.2, multiplicative speckle noise of mean 0 and variance 0.2, and impulse noise of density 0.2 to the brain proton density slice image I .

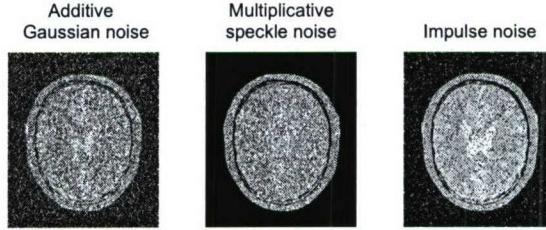


FIGURE 2. An illustration of the addition of various types of noise to the image I .

In this paper, we will study the problem of image registration in the presence of high levels of impulse noise. We will determine the impulse noise density level at which ordinary registration methods fail, and we will present a multiscale registration algorithm that enables accurate registration for noise levels higher than those at which ordinary methods fail. To study the effect of varying noise densities on the registration process, we add impulse noise of increasing densities δ to the brain proton density slice image I , and register the (non-noisy) translated image T with the noisy images. Let I_δ denote the image I with added impulse noise of density δ . In Figure 3, we illustrate the noisy images I_δ for increasing values of δ .

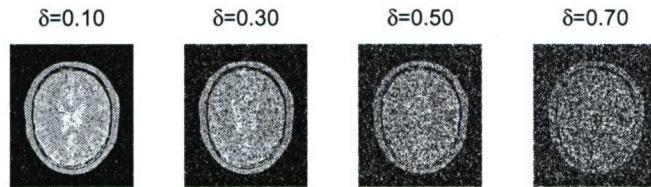


FIGURE 3. An illustration of adding impulse noise of increasing densities δ to the image I .

REMARK. Although in this paper we present the results of image registration experiments using only impulse noise, we have also conducted numerous experiments using other types of noise, including additive Gaussian noise and speckle noise. The results obtained with all other types of noise are similar.

3.2. Registration results. For increasing noise densities δ , we register T with I_δ using various registration methods. Recall that the image T is the result of translating the original image I 13 units in X and 17 units in Y , and that I_δ is the result of adding uniform impulse noise of density δ to the image I . Since T is a rigid transformation of I , we will restrict the registration process to linear transformations; i.e., we will consider optimal linear registrations. The optimal transformation ϕ produced by the optimal linear registration process will consist of two parameters, namely X - and Y -translation values. We will let ϕ_X and ϕ_Y denote the X - and Y -translation parameters, respectively, of the optimal transformation ϕ . For comparison purposes, we will perform the optimal linear registration using the mean squares, normalized correlation, and mutual information metrics.

We use the following parameters for the registration algorithms. For the mean squares and normalized correlation registration algorithms, we use the regular step gradient descent optimizer. Due to the stochastic nature of the metric computation in the mutual information algorithm, the regular step gradient descent optimizer does not work well in the case of mutual information. Instead, we use the gradient descent optimizer with a user-specified learning rate of 20.0. See [7] for a detailed

TABLE 1. The results obtained upon registering the translated image T with the noisy image I_δ , where δ is the impulse noise density; ϕ_X and ϕ_Y denote the X - and Y -translation values of the optimal transformation ϕ produced by the registration algorithm, and n is the number of iterations until convergence. The actual translation values are 13 units in X and 17 units in Y .

δ	Mean Squares			Normalized Correlation			Mutual Information		
	ϕ_X	ϕ_Y	n	ϕ_X	ϕ_Y	n	ϕ_X	ϕ_Y	n
0.00	12.99	17.00	18	13.01	17.00	18	12.75	17.03	200
0.10	12.99	17.01	28	12.99	17.01	20	12.83	16.88	200
0.20	13.03	16.98	17	13.04	16.98	19	12.98	16.64	200
0.30	12.97	17.03	28	13.02	17.02	11	13.02	17.02	200
0.40	18.89	7.16	15	8.05	1.30	13	11.08	9.72	200
0.50	2.16	7.06	19	9.09	2.18	8	9.72	7.12	200
0.60	29.81	3.19	40	4.08	0.24	7	4.57	5.17	200
0.70	2.08	1.14	13	3.11	2.13	12	3.08	2.86	200

discussion of these parameters. Finally, we set the maximum number of iterations for each algorithm to 200. As we shall see, mean squares and normalized correlation registrations typically converge very quickly to the optimal value. Mutual information, on the other hand, often does not actually reach the true optimal solution but instead oscillates within one or two pixels of the optimal solution (generally after 100-150 iterations). By reducing the learning rate, we can increase the likelihood of convergence, but this increases the computation time significantly without improving the accuracy of the solution.

For each of these three registration algorithms, and for each δ we record the X - and Y -translation parameters, denoted by ϕ_X and ϕ_Y , respectively, of the optimal transformation ϕ produced by the registration process. We also record the number of iterations n until convergence. The results are shown in Table 1. Recall that the actual translation values are 13 units in X and 17 units in Y . We also record the number of iterations until convergence, which we denote by n .

The results presented in Table 1 indicate that optimal linear registration in the presence of impulse noise fails when the impulse noise density in the fixed image reaches approximately 0.40, regardless of the metric used.

4. Denoising.

4.1. Denoising techniques. In this section, we discuss various denoising techniques. Image denoising is a fundamental problem in image processing, and there has been much research and progress on the subject. As our primary interest is not denoising but the problem of image registration of noisy images, we do not focus on the general problem of image denoising. Instead, we present a few of the most common and computationally simple denoising techniques. We will then apply these techniques to one of our noisy images and study the effect of denoising on the image registration techniques. In particular, in Section 3, we saw that ordinary optimal linear registration of noisy images failed when the impulse noise density was greater than 0.40. Also in this section, we shall determine whether or not denoising prior

to registration enables successful registration of noisy images for which registration failed previously.

Spatial filtering is the traditional approach to removing noise from images. Spatial filters use the assumption that noise occupies the higher regions of the frequency spectrum, and thus they attenuate high spatial frequencies. Local spatial filtering is a process in which the value of a given pixel in the filtered image is computed by applying some algorithm to the pixel values in a neighborhood of the given pixel. Typical implementations of spatial filters include mean filtering, median filtering, and Gaussian smoothing. Mean filtering computes the value of each output pixel by computing the statistical mean of the neighborhood of the corresponding input pixel. Thus, applying a mean filter to a noisy image reduces the amount of variation in gray-level intensity between pixels. Although this filter is computationally easy to implement, it is sensitive to the presence of outliers. Median filtering, which computes the value of each output pixel by computing the statistical median of the neighborhood of the corresponding input pixel, is more robust to the presence of outliers, and is thus commonly used for removing impulse noise from images. Convolution with a Gaussian kernel is another commonly used spatial filtering technique. See [23] for an overview of classical spatial filtering techniques.

In Figure 4, we illustrate the effect of applying a mean, median, and Gaussian convolution filter to the noisy image $I_{0.70}$, the brain proton density slice image with impulse noise of density 0.70. As is indicated in Figure 4, spatial filters smooth the data to remove noise but also blur edges.

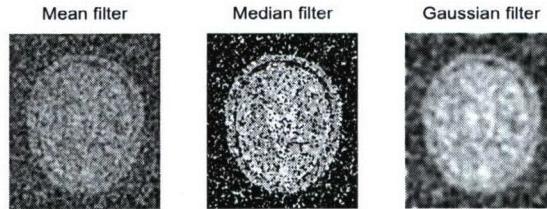


FIGURE 4. The results of applying mean, median, and Gaussian filters to the brain proton density slice image with impulse noise of density 0.70.

More advanced denoising techniques that remove noise more effectively while preserving edges include wavelet-based methods [10], total variation methods [19], and PDE-based anisotropic diffusion methods [13], to name a few. Total variation denoising reduces the total variation of the image, and thus removes noise, textures, and fine-scale details while preserving edges. In Figure 5, we illustrate the effect of applying these denoising techniques to the noisy image $I_{0.70}$, the brain proton density slice image with impulse noise of density 0.70.

4.2. Registration results after denoising. In this section, we register the translated image T with the denoised images illustrated in Figures 4 and 5. As in Section

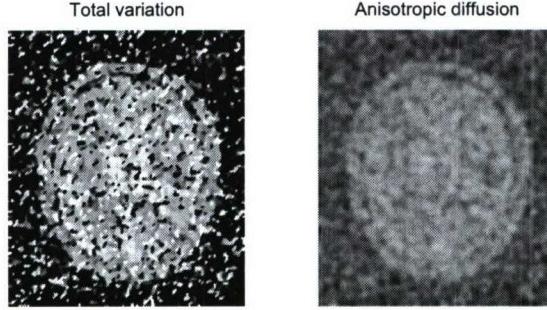


FIGURE 5. The results of applying the Osher-Rudin total variation and the Perona-Malik anisotropic diffusion denoising algorithms to the brain proton density slice image with impulse noise of density 0.70.

TABLE 2. The results obtained upon registering the translated image T with the denoised images obtained upon applying median, mean, and Gaussian convolution filters to the noisy image $I_{0.70}$. ϕ_X and ϕ_Y are the X - and Y -translation values of the optimal transformation ϕ produced by the registration algorithm, n is the number of iterations until convergence. The actual translation values are 13 units in X and 17 units in Y .

Denoising Technique	Mean Squares			Normalized Correlation			Mutual Information		
	ϕ_X	ϕ_Y	n	ϕ_X	ϕ_Y	n	ϕ_X	ϕ_Y	n
Mean Filtering	31.83	1.15	46	16.88	1.11	29	5.39	5.30	200
Median Filtering	18.87	1.26	31	2.38	6.90	34	4.39	4.06	200
Gaussian Filtering	18.86	-0.76	31	2.19	0.25	11	7.38	7.37	200
Total Variation	6.11	4.26	19	5.29	9.15	15	6.30	6.23	15
Anisotropic Diffusion	2.10	1.13	11	4.09	6.22	10	10.62	14.77	200

3, we use mean squares, normalized correlation, and mutual information optimal linear registration. For each registration method, we let ϕ denote the optimal transformation produced by the registration algorithm, and we let ϕ_X and ϕ_Y the X - and Y -translation parameters of the optimal transformation ϕ . We denote by n the number of iterations of each registration algorithm until convergence. We record the results in Table 2. The moving image in each case is the translated image T ; recall that the actual translation values are 13 in X and 17 in Y .

The results presented in Table 2 indicate that the application of some of the classical as well as modern denoising techniques prior to registration does not enable successful registration of the noisy image $I_{0.70}$ with the translated image T . Although the more advanced denoising techniques such as total variation and anisotropic diffusion result in translation values that are closer to the actual values, particularly when mutual information registration is used, we conclude from Table 2 that denoising prior to registration does not produce accurate registration results for images that contain high levels of noise.

5. Multiscale decomposition. In this section, we present the multiscale image representation using hierarchical (BV, L^2) decompositions of [20]. The multiscale decomposition will provide a hierarchical expansion of an image that separates the essential features of the image (such as large shapes and edges) from the fine scales of the image (such as details and noise). The decomposition is hierarchical in the sense that it will produce a series of expansions of the image that resolve increasingly finer scales, and hence will include increasing levels of detail. We will eventually apply the multiscale decomposition algorithm to the problem of image registration in the presence of noise, and will demonstrate the accuracy of the multiscale registration technique for noisy images such as those that were considered in Sections 3 and 4.

We will use the following mathematical spaces in the decomposition algorithm. The space of functions of bounded variation, BV , is defined by

$$BV = \left\{ f \quad \left| \quad \|f\|_{BV} := \sup_{h \neq 0} |h|^{-1} \|f(\cdot + h) - f(\cdot)\|_{L^1} < \infty \right. \right\}.$$

We will also use the Sobolev space $W^{-1,\infty}$ with norm given by:

$$\|f\|_{W^{-1,\infty}} := \sup_g \left[\int \frac{f(x)g(x)}{\|g\|_{W^{1,1}}} dx \right],$$

where $\|g\|_{W^{1,1}} := \|\nabla g\|_{L^1}$.

5.1. The hierarchical decomposition. Define the J -functional $J(f, \lambda)$ as follows:

$$J(f, \lambda) := \inf_{u+v=f} \lambda \|v\|_{L^2}^2 + \|u\|_{BV}, \quad (18)$$

where $\lambda > 0$ is a scaling parameter that separates the L^2 and BV terms. This functional $J(f, \lambda)$ was introduced in the context of image processing by Rudin, Osher, and Fatemi [19]. They suggested the following. Let $[u_\lambda, v_\lambda]$ denote the minimizer of $J(f, \lambda)$. The BV component, u_λ , captures the coarse features of the image f , while the L^2 component, v_λ , captures the finer features of f such as noise. This model is effective in denoising images while preserving edges, though it requires prior knowledge of the noise scaling λ .

Tadmor, et al. propose in [20] an alternative point of view in which the minimization of $J(f, \lambda)$ is interpreted as a decomposition $f = u_\lambda + v_\lambda$, where u_λ extracts the edges of f and v_λ extracts the textures of f . This interpretation depends on the scale λ , since texture at scale λ consists of edges when viewed under a refined scale (2^λ , for example). We refer to $v_\lambda = f - u_\lambda$ as the residual of the decomposition. Upon decomposing $f = u_\lambda + v_\lambda$, we proceed to decompose v_λ as follows:

$$v_\lambda = u_{2\lambda} + v_{2\lambda},$$

where

$$[u_{2\lambda}, v_{2\lambda}] = \operatorname{arginf}_{u+v=v_\lambda} J(v_\lambda, 2\lambda).$$

Thus, we obtain a two-scale representation of f given by $f \cong u_\lambda + u_{2\lambda}$, where now $v_{2\lambda} = f - (u_\lambda + u_{2\lambda})$ is the residual. Next we decompose $v_{2\lambda}$ and continue this process, which results in the following hierarchical multiscale decomposition of f . Starting with an initial scale $\lambda = \lambda_0$, we obtain an initial decomposition of the image f :

$$f = u_0 + v_0, \quad [u_0, v_0] = \operatorname{arginf}_{u+v=f} J(f, \lambda_0).$$

We then refine this decomposition to obtain

$$v_j = u_{j+1} + v_{j+1}, \quad [u_{j+1}, v_{j+1}] = \operatorname{arginf}_{u+v=v_j} J(v_j, \lambda_0 2^{j+1}), \quad j = 0, 1, \dots$$

After k steps of this process, we have:

$$f = u_0 + v_0 = u_0 + u_1 + v_1 = u_0 + u_1 + u_2 + v_2 = \dots = u_0 + u_1 + \dots + u_k + v_k,$$

which is a multiscale image decomposition $f \sim u_0 + u_1 + \dots + u_k$, with a residual v_k . As k increases, the u_k components resolve edges with increasing scales $\lambda_k = \lambda_0 2^k$.

5.2. Implementation.

5.2.1. Initialization. As described in [20], the initial scale λ_0 should capture the smallest oscillatory scale in f , given by

$$\frac{1}{2\lambda_0} \leq \|f\|_{W^{-1,\infty}} \leq \frac{1}{\lambda_0}. \quad (19)$$

However, in practice, we may not be able to determine the size of $\|f\|_{W^{-1,\infty}}$, so we determine the initial choice of λ_0 experimentally. Following [20], for the applications presented in this paper, we will use $\lambda_0 = 0.01$ and $\lambda_j = \lambda_0 2^j$.

5.2.2. Numerical discretization. We follow the numerical algorithm of [20] for the construction of our hierarchical decomposition. In each step, we use finite-difference discretization of the Euler-Lagrange equations associated with the $J(v_j, \lambda_{j+1})$ to obtain the next term, u_{j+1} , in the decomposition of the image f . The Euler-Lagrange equation associated with the minimization of the functional $J(f, \lambda)$ given in equation (18) is

$$u_\lambda - \frac{1}{2\lambda} \operatorname{div} \left(\frac{\nabla u_\lambda}{|\nabla u_\lambda|} \right) = f,$$

with the Neumann boundary conditions:

$$\frac{\partial u_\lambda}{\partial n} \Big|_{\partial\Omega} = 0, \quad (20)$$

where $\partial\Omega$ is the boundary of the domain Ω and n is the unit outward normal.

We thus obtain an expansion $f \sim \sum_{j=0}^k u_j$, where the u_j are constructed as approximate solutions of the recursive relation given by the following elliptic PDE:

$$u_{j+1} - \frac{1}{2\lambda_{j+1}} \operatorname{div} \left(\frac{\nabla u_{j+1}}{|\nabla u_{j+1}|} \right) = -\frac{1}{2\lambda_j} \operatorname{div} \left(\frac{\nabla u_j}{|\nabla u_j|} \right). \quad (21)$$

Note that $J(f, \lambda)$ contains a singularity when $|\nabla u_\lambda| = 0$. To remove this singularity, we replace $J(f, \lambda)$ by the regularized functional

$$J^\epsilon(f, \lambda) := \inf_{u+v=f} \left\{ \lambda ||v||_{L^2}^2 + \int_{\Omega} \sqrt{\epsilon^2 + |\nabla u|^2} dx dy \right\}, \quad (22)$$

and at each step, we find the minimizer u_λ of J^ϵ . The Euler-Lagrange equation for the regularized J^ϵ functional is

$$u_\lambda - \frac{1}{2\lambda} \operatorname{div} \left(\frac{\nabla u_\lambda}{\sqrt{\epsilon^2 + |\nabla u_\lambda|^2}} \right) = f \in \Omega,$$

with Neumann boundary conditions.

To numerically implement the method, we cover the domain Ω with a grid ($x_i := ih, y_j := jh$), and discretize the elliptic PDE of equation (21) as follows:

$$\begin{aligned} u_{i,j} &= f_{i,j} + \frac{1}{2\lambda} D_{-x} \left[\frac{1}{\sqrt{\epsilon^2 + (D_{+x} u_{i,j})^2 + (D_{0y} u_{i,j})^2}} D_{+x} u_{i,j} \right] \\ &\quad + \frac{1}{2\lambda} D_{-y} \left[\frac{1}{\sqrt{\epsilon^2 + (D_{0x} u_{i,j})^2 + (D_{+y} u_{i,j})^2}} D_{+y} u_{i,j} \right] \\ &= f_{i,j} + \frac{1}{2h^2} \left[\frac{u_{i+1,j} - u_{i,j}}{\sqrt{\epsilon^2 + (D_{+x} u_{i,j})^2 + (D_{0y} u_{i,j})^2}} \right. \\ &\quad \left. - \frac{u_{i,j} - u_{i-1,j}}{\sqrt{\epsilon^2 + (D_{-x} u_{i,j})^2 + (D_{0y} u_{i-1,j})^2}} \right] \\ &\quad + \frac{1}{2h^2} \left[\frac{u_{i,j+1} - u_{i,j}}{\sqrt{\epsilon^2 + (D_{0x} u_{i,j})^2 + (D_{+y} u_{i,j})^2}} \right. \\ &\quad \left. - \frac{u_{i,j} - u_{i,j-1}}{\sqrt{\epsilon^2 + (D_{0x} u_{i,j-1})^2 + (D_{-y} u_{i,j})^2}} \right], \end{aligned} \quad (23)$$

where D_+ , D_- , and D_0 denote the forward, backward, and centered divided differences, respectively. To solve the discrete regularized Euler-Lagrange equations (24), we use the Gauss-Siedel iterative method to obtain:

$$\begin{aligned}
u_{i,j}^{n+1} = f_{i,j} + \frac{1}{2h^2} & \left[\frac{u_{i+1,j}^n - u_{i,j}^{n+1}}{\sqrt{\epsilon^2 + (D_{+x}u_{i,j}^n)^2 + (D_{0y}u_{i,j}^n)^2}} \right. \\
& - \frac{u_{i,j}^{n+1} - u_{i-1,j}^n}{\sqrt{\epsilon^2 + (D_{-x}u_{i,j}^n)^2 + (D_{0y}u_{i-1,j}^n)^2}} \Big] \\
& + \frac{1}{2h^2} \left[\frac{u_{i,j+1}^n - u_{i,j}^{n+1}}{\sqrt{\epsilon^2 + (D_{0x}u_{i,j}^n)^2 + (D_{+y}u_{i,j}^n)^2}} \right. \\
& \left. - \frac{u_{i,j}^{n+1} - u_{i,j-1}^n}{\sqrt{\epsilon^2 + (D_{0x}u_{i,j-1}^n)^2 + (D_{-y}u_{i,j}^n)^2}} \right]. \tag{24}
\end{aligned}$$

To satisfy the Neumann boundary conditions (20), we first reflect f outside Ω by adding grid lines on all sides of Ω . As the initial condition, we set $u_{i,j}^0 = f_{i,j}$. We iterate this numerical scheme for $n = 0, 1, \dots, N$ until $\|u^{n_\infty} - u^{n_\infty-1}\|$ is less than some preassigned value so that $u_{i,j}^{n_\infty}$ is an accurate approximation of the fixed point steady solution u_λ .

Finally, we denote the final solution $u_\lambda := \{u_{i,j}^{n_\infty}\}_{i,j}$. To obtain the hierarchical multiscale decomposition, we reiterate this process, each time updating f and λ in the following way:

$$\begin{aligned}
f_{new} & \leftarrow f_{current} - u_\lambda, \\
\lambda_{new} & \leftarrow 2\lambda_{current}. \tag{25}
\end{aligned}$$

That is, at each step, we apply the $J(f_{current} - u_\lambda, 2\lambda)$ minimization to the residual $f_{current} - u_\lambda$ of the previous step. Taking $\lambda_j = \lambda_0 2^j$, we obtain after k steps a hierarchical multiscale decomposition $f = u_{\lambda_0} + u_{\lambda_1} + \dots + u_{\lambda_k} + v_{\lambda_k}$, where we write $u_{\lambda_j} = u_j$. We call the u_j , $j = 1, 2, \dots, k$ the components of f and the v_k the residuals.

EXAMPLE 1. Decomposition of a noisy image. We apply the hierarchical multiscale decomposition of [20] as described in Section 5 to the noisy image $I_{0.70}$ in Figure 3, using the following parameters: $m = 12$ hierarchical steps, $\lambda_0 = 0.01$, $\lambda_j = \lambda_0 2^j$, $\epsilon = 0.001$, $n = 10$, and $h = 1$. In Figures 6 and 7, we illustrate the components u_{λ_j} and the residuals v_{λ_j} for this decomposition. Note that in each hierarchical step, an additional amount of texture is seen in the components. Further, the noise is not seen in the first few components, while most of the texture is kept, and the noise only reappears as the refined scales reach the same scales as the noise itself. Our goal is to use this multiscale decomposition to register the noisy image $I_{0.70}$ with the translated image T .

6. Multiscale registration. Consider again the noisy images I_δ shown in Figure 3 with impulse noise of increasing densities δ . Recall that in Section 3, we demonstrated that registration of the the translated image T with the noisy image I_δ

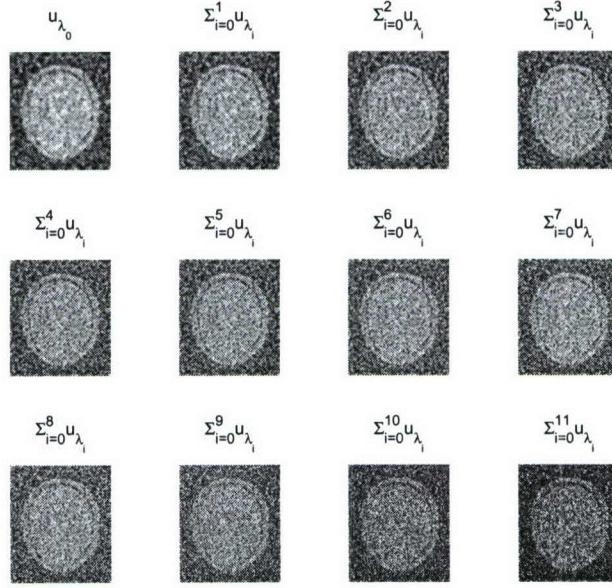


FIGURE 6. Multiscale decomposition of the noisy image $I_{0.70}$ shown in Figure 3.

failed when $\delta \geq 0.40$, regardless of the metric used in the optimal linear registration process. Moreover, registration using these classical methods failed even after denoising the noisy image using various standard denoising techniques, as demonstrated in Section 4. In this section, we present new methods for image registration that allow for a successful registration of the translated image T with the noisy images I_δ for values of the noise density δ significantly greater than the levels at which classical registration and registration after denoising fail. These registration techniques will be based on the hierarchical multiscale decomposition described in Section 5.

Consider two images A and B , and suppose that we want to register image B with image A . Suppose that one or both of the images contains a significant amount of noise. If only one of the images is noisy, we assume that it is image A . We propose the following multiscale registration method. First, we apply the multiscale hierarchical decomposition to both images. Let m denote the number of hierarchical steps used for the multiscale decompositions. For ease of notation, given an image f , we let

$$C_k(f) := \sum_{i=0}^k u_{\lambda_k} \quad (26)$$

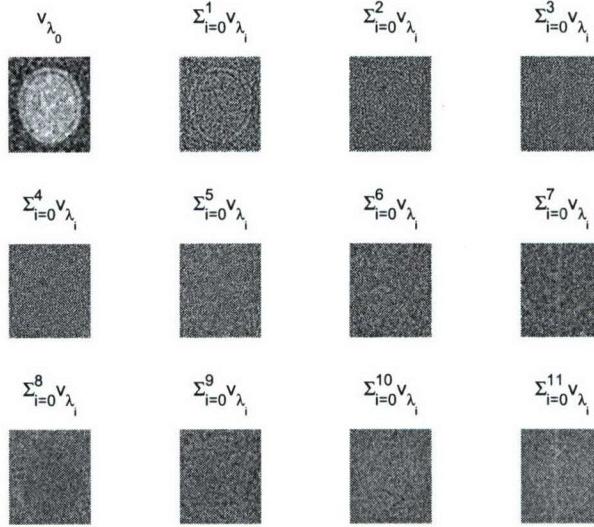


FIGURE 7. The residuals of the multiscale decomposition of the noisy image $I_{0.70}$ shown in Figure 3.

denote the k^{th} component of the image f , $k = 0, 1, \dots, m - 1$, obtained as in Section 5. Thus $C_k(A)$ will denote the k^{th} component of the image A , and $C_k(B)$ will denote the k^{th} component of image B .

We will present two algorithms; in the first, we register image B with the components of image A , and in the second, we register the components of image B with the components of image A .

6.1. Algorithm I: One-node multiscale registration. In our first multiscale registration algorithm, we register image B with the k^{th} component of A , for $k = 0, 1, \dots, m - 1$. This is illustrated by the schematic in Figure 8.

We refer to this algorithm as a *one-node multiscale registration algorithm* because in each of the m registrations prescribed by the algorithm, the moving image is always the image B . We only use the multiscale components of the fixed image A for the one-node algorithm.

Let ϕ_k denote the optimal transformation produced by the registration algorithm upon registering B with $C_k(A)$, $k = 0, 1, \dots, m - 1$. Recall that $C_0(A)$ contains only the coarsest scales of the image A , and as k increases, $C_k(A)$ contains increasing levels of detail (and hence, noise) of the image A . Thus, we expect that registration of image B with $C_k(A)$ should give an improvement compared to ordinary registration for the first few values of k . As k increases, however, we expect that eventually the component $C_k(A)$ will become too noisy to give successful registration.

Upon determining the transformations ϕ_k with a suitable registration algorithm (e.g., an optimal linear registration), we have several options for defining the optimal transformation Φ that should bring the image B into spatial alignment with the

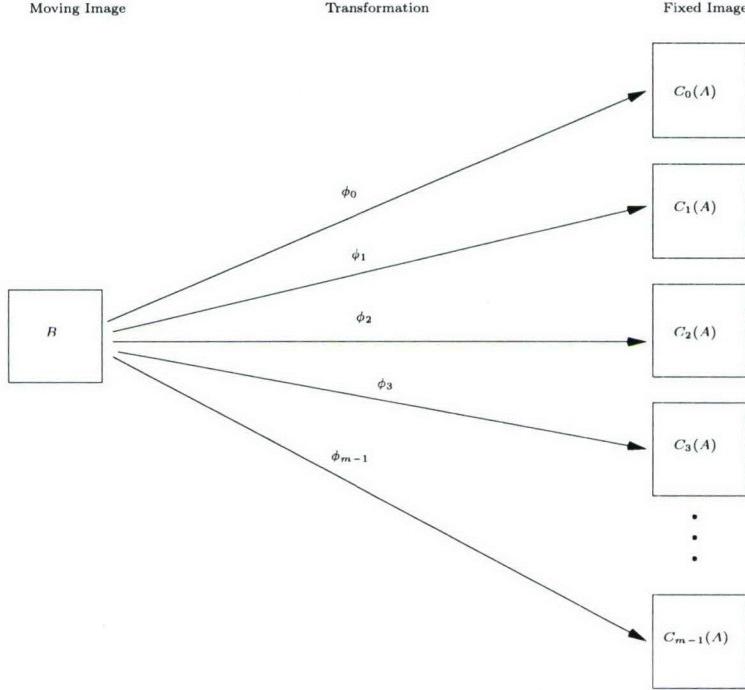


FIGURE 8. This schematic represents a one-node multiscale image registration algorithm in which we register the moving image B with the k^{th} component of the fixed image A , for $k = 0, 1, \dots, m-1$, where m is the number of hierarchical steps used for the multiscale decompositions.

image A . The first option would be to take into account the registration parameters corresponding to the coarse scales only, i.e., the first few values of k , for which we expect a more accurate registration. Upon determining the number of registrations that we wish to take into account, we could then estimate Φ by averaging the registration parameters corresponding to those coarse scale registrations. A second option would be to define Φ as a weighted average of the ϕ_k ; i.e.,

$$\Phi := \frac{1}{m} \sum_{k=0}^{m-1} a_k \phi_k, \quad (27)$$

where the weights a_k are appropriately chosen non-negative real numbers such that $\sum a_k = m$. For example, we could perform a statistical analysis on the registration parameters corresponding to the ϕ_k , and use the mean and standard deviation (or the mean and standard deviation of the first several values) to determine the weights a_k .

6.2. Algorithm II: Multi-node multiscale registration. In our second multiscale registration algorithm, we register the k^{th} component of image B with the

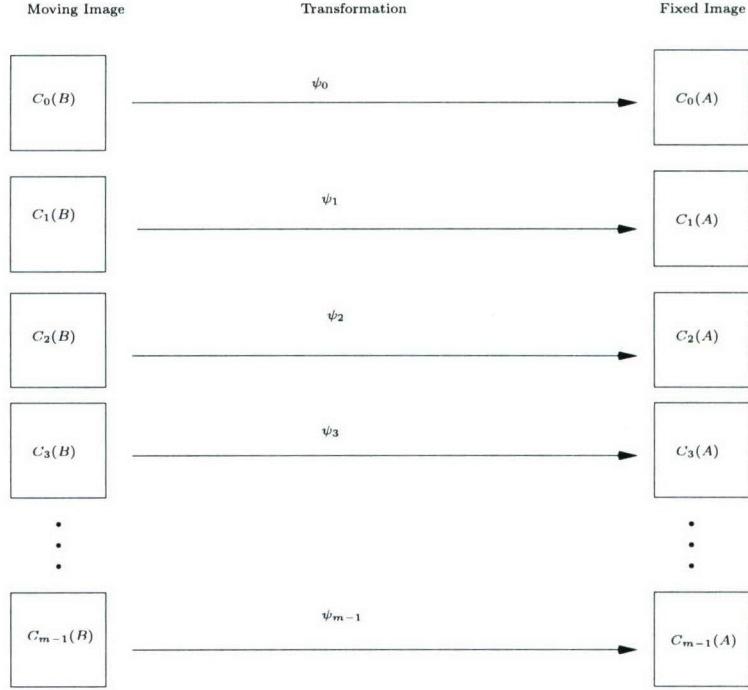


FIGURE 9. This schematic represents a multi-node multiscale image registration algorithm in which we register the k^{th} component of the moving image B with the k^{th} component of the fixed image A , for $k = 0, 1, \dots, m - 1$, where m is the number of hierarchical steps used for the multiscale decompositions.

k^{th} component of image A , for $k = 0, 1, 2, \dots, m - 1$, as illustrated by the schematic in Figure 9.

We refer to this algorithm as a *multi-node multiscale registration algorithm* because in each of the m registrations prescribed by the algorithm, we consider both the components of the fixed image A and the components of the moving image B .

Let ψ_k denote the optimal transformation produced by the registration algorithm upon registering $C_k(B)$ with $C_k(A)$, $k = 0, 1, \dots, m - 1$. As before, we expect that registration of $C_k(B)$ with $C_k(A)$ should give an improvement compared to ordinary registration for the first few values of k . As k increases, however, we expect that eventually the components $C_k(A)$ and $C_k(B)$ will become too noisy to register successfully. Since this algorithm considers components of both images, we expect that it will be particularly successful in the case in which both images are noisy.

As in the case of the one-node multiscale registration algorithm, we can define the optimal transformation Ψ that should bring image B into spatial alignment with image A either by taking into account only the first few registration results corresponding to registration of the coarse scales, and averaging the registration parameters corresponding to these first few registrations, or by computing a weighted average:

TABLE 3. The registration results upon registering T with $I_{0.70}$ using Algorithm I. Here, we use $m = 12$ hierarchical steps to decompose the noisy image, so we perform $m = 12$ registration simulations. The transformation parameters ϕ_X and ϕ_Y are the X - and Y -translation parameters of the optimal transformation ϕ produced by the registration algorithm. The actual translation values are 13 in X and 17 in Y . The moving image in all simulations is the translated image T .

Fixed Image	Mean Squares			Normalized Correlation			Mutual Information		
	ϕ_X	ϕ_Y	n	ϕ_X	ϕ_Y	n	ϕ_X	ϕ_Y	n
$I_{0.70}$	4.57	5.17	200	2.08	1.14	7	4.08	0.24	7
$C_0(I_{0.70})$	12.65	16.36	200	3.08	1.11	12	3.11	0.17	9
$C_1(I_{0.70})$	12.69	16.78	200	2.08	3.08	14	2.13	2.12	12
$C_2(I_{0.70})$	12.56	16.79	200	2.11	3.08	14	2.14	3.11	15
$C_3(I_{0.70})$	12.53	16.76	200	3.08	2.11	14	3.11	2.14	7
$C_4(I_{0.70})$	12.48	16.76	200	24.88	1.16	36	18.86	1.18	30
$C_5(I_{0.70})$	12.46	16.78	200	40.80	1.07	52	0.21	1.18	11
$C_6(I_{0.70})$	12.43	16.80	200	28.86	0.15	46	27.84	2.19	42
$C_7(I_{0.70})$	12.43	16.79	200	-2.87	4.11	15	0.18	3.14	12
$C_8(I_{0.70})$	12.43	16.74	200	25.89	3.12	40	-1.84	4.12	14
$C_9(I_{0.70})$	9.33	9.41	200	6.05	4.12	12	7.99	2.08	16
$C_{10}(I_{0.70})$	8.44	8.32	200	-3.92	8.12	21	4.09	3.15	16
$C_{11}(I_{0.70})$	6.96	6.46	200	8.97	6.13	13	3.65	1.17	27

$$\Psi := \frac{1}{m} \sum_{k=0}^{m-1} b_k \psi_k, \quad (28)$$

where the weights b_k are appropriately chosen non-negative real numbers such that $\sum b_k = m$.

7. Examples of multiscale registration.

7.1. A noisy fixed image. In this section, we use the multiscale registration technique described in Section 6 to register the translated (non-noisy) image T with the noisy image $I_{0.70}$. Recall that $I_{0.70}$ is the image obtained upon adding impulse noise of density 0.70 to the brain proton density slice image I . As before, let $C_k(I_{0.70})$ denote the k^{th} component in the multiscale decomposition of $I_{0.70}$, for $k = 0, 1, \dots, m$, obtained as in Section 6. We perform the multiscale decomposition using $m = 12$ hierarchical steps, $\lambda_0 = 0.01$, and $\lambda_j = \lambda_0 2^j$. In Table 3, we present the results of $m = 12$ registration simulations, obtained upon registering T with $C_k(I_{0.70})$, $k = 0, 1, \dots, 11$, using Algorithm I of Section 6.1. For each registration, we let ϕ denote the optimal transformation produced by the registration algorithm, and we let ϕ_X and ϕ_Y the X - and Y -translation parameters of the optimal transformation ϕ . The moving image in each registration is the translated image T . For reference, we also include in the first line of Table 3 the parameters obtained using ordinary registration.

It is clear from the results presented in Table 3 that the results obtained using mean squares and normalized correlation methods are completely inaccurate. Thus, the one-node multiscale algorithm did not produce meaningful results for these metrics. For mutual information, however, the X and Y translation parameters are clustered around 12.5 units in X and 16.8 units in Y for $k = 0, 1, \dots, 8$, but then are significantly different for the remaining values of k . We expected that the multiscale registration results would be an accurate approximation of the actual transformation Φ for small values of k , but then would deviate as k became sufficiently large, because as k becomes large, increasing scales of detail (and hence, noise) appear in the component C_k . Thus, even without knowing the actual values of the X - and Y -translations, it makes sense to take into account only the parameters corresponding to the first 9 registrations ($k = 0, 1, \dots, 8$). Averaging the translation parameters for the first 9 registrations, we obtain $\Phi_X = 12.52$ and $\Phi_Y = 16.73$. Since the actual values are 13 in X and 17 in Y , we see that multiscale mutual information registration produced very accurate results in this case, and indeed is a significant improvement compared to ordinary registration as well as to classical and modern denoising followed by registration.

Next, we provide the results obtained with Algorithm II by registering the multiscale components of the translated image T with the multiscale components of the noisy image $I_{0.70}$. Let $C_k(T)$ and $C_k(I_{0.70})$ denote the multiscale components of T and $I_{0.70}$, respectively, obtained through the multiscale decomposition presented in Section 5. As before, we use $m = 12$ hierarchical steps, $\lambda_0 = 0.01$, and $\lambda_j = \lambda_0 2^j$ to perform the decomposition. In Table 4, we present the results of $m = 12$ registration simulations, obtained upon registering $C_k(T)$ with $C_k(I_{0.70})$, $k = 0, 1, \dots, 11$. For each registration, we let ψ denote the optimal transformation produced by the registration algorithm, and let ψ_X and ψ_Y denote the X - and Y -translation parameters of the optimal transformation ψ . For reference, we also include in the first line of Table 4 the parameters obtained using ordinary registration.

To estimate the transform parameters Ψ_X and Ψ_Y , we note that for mutual information, the translation parameters ψ_x and ψ_y are clustered together for the first 9 registrations, and for mean squares and normalized correlation, the values are clustered together for the first 2 registrations. Thus for mutual information we determine Ψ by averaging the parameters corresponding to the first 9 registrations, and for mean squares and normalized correlation, we average the first 2 values. In Table 5, we present the X - and Y -translation values corresponding to these averages.

REMARK. Since the actual translation values are 13 in X and 17 in Y , we see that the multinode multiscale registration of the translated image T with the noisy image $I_{0.70}$ produces very accurate results for each of the three optimal linear registration metrics considered here (mean squares, normalized correlation, and mutual information). The main difference between the results obtained with Algorithm I and Algorithm II is the accurate registration of the coarse scales obtained with Algorithm II.

7.2. Noisy fixed and moving images. In this section, we consider the registration problem in which both the fixed and moving images are noisy. Consider the noisy images $I_{0.40}$ and $T_{0.40}$, where T , as before, is the result of translating I 13 units in X and 17 units in Y , and A_δ denotes the image obtained by adding impulse noise of density δ to the image A . The noisy images are shown in Figure 10.

TABLE 4. The registration results obtained with Algorithm II. Here we register the k^{th} multiscale component $C_k(T)$ of the translated image T with the k^{th} multiscale component $C_k(I_{0.70})$ of the noisy image $I_{0.70}$ obtained via the multiscale decomposition discussed in Section 5. Here, we use $m = 12$ hierarchical steps to decompose the noisy image, so we perform $m = 12$ registration simulations. The transformation parameters ψ_X and ψ_Y are the X - and Y -translation parameters of the optimal transformation ψ produced by the registration algorithm. The actual translation values are 13 in X and 17 in Y .

Fixed and Moving Images	Mean Squares		Normalized Correlation		Mutual Information	
	ϕ_X	ϕ_Y	ϕ_X	ϕ_Y	ϕ_X	ϕ_Y
$I_{0.70}$ and T	4.57	5.18	2.08	1.14	4.08	0.24
$C_0(I_{0.70})$ and $C_0(T)$	12.69	16.66	12.29	17.72	12.96	17.08
$C_1(I_{0.70})$ and $C_1(T)$	12.67	16.87	13.70	17.75	12.99	17.67
$C_2(I_{0.70})$ and $C_2(T)$	12.59	16.86	20.77	5.20	16.84	4.31
$C_3(I_{0.70})$ and $C_3(T)$	12.55	16.82	3.19	0.31	4.20	4.23
$C_4(I_{0.70})$ and $C_4(T)$	12.52	16.83	2.20	2.24	26.74	5.18
$C_5(I_{0.70})$ and $C_5(T)$	12.51	16.84	31.65	2.23	14.90	6.27
$C_6(I_{0.70})$ and $C_6(T)$	12.49	16.87	30.69	6.16	19.87	4.29
$C_7(I_{0.70})$ and $C_7(T)$	12.48	16.85	33.64	3.16	29.64	3.32
$C_8(I_{0.70})$ and $C_8(T)$	12.53	16.71	28.81	3.22	1.26	1.29
$C_9(I_{0.70})$ and $C_9(T)$	9.26	9.36	2.13	3.13	17.93	3.21
$C_{10}(I_{0.70})$ and $C_{10}(T)$	8.80	8.61	2.12	3.12	32.63	3.14
$C_{11}(I_{0.70})$ and $C_{11}(T)$	6.95	6.34	34.74	2.10	4.13	5.08

TABLE 5. The translation parameters Ψ_X and Ψ_Y obtained by averaging the parameters corresponding to the coarse scale registrations. The actual translation values are 13 in X and 17 in Y .

	Mean Squares	Normalized Correlation	Mutual Information
Ψ_X	12.56	12.99	12.98
Ψ_Y	16.82	17.74	17.37

Before applying our multiscale registration algorithm, we attempt to register $T_{0.40}$ with $I_{0.40}$ using the three registration methods mean squares, normalized correlation, and mutual information. The results shown in Table 6 indicate that registration of the noisy images fails, regardless of the metric used in the optimal linear registration algorithm.

Since ordinary registration of the noisy images fails, we register the images using Algorithm II, the multi-node multiscale registration technique. First, we perform the multiscale decomposition discussed in Section 5 to both noisy images, again using $m = 12$ hierarchical steps, initial scale $\lambda_0 = 0.01$, and $\lambda_j = 2^j \lambda_0$. Let $C_k(I_{0.40})$ and $C_k(T_{0.40})$ denote the k^{th} component in the multiscale decomposition of $I_{0.40}$

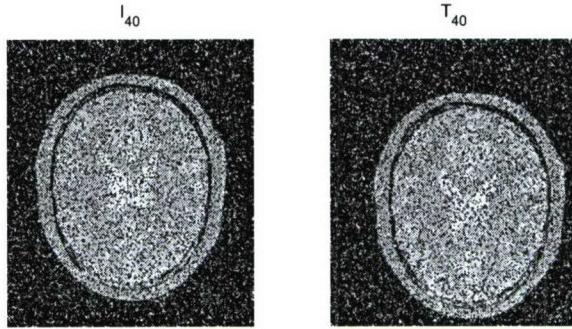


FIGURE 10. The original image I and translated image T with impulse noise of density $\delta = 0.40$

TABLE 6. The results of registering the noisy translated image $T_{0.40}$ with the noisy image $I_{0.40}$, using three different metrics. The actual translation values are 13 in X and 17 in Y .

Registration Method	ϕ_X	ϕ_Y
Mean Squares	11.02	7.04
Normalized Correlation	3.05	0.99
Mutual Information	5.03	2.54

and $T_{0.40}$, respectively. Since both images are noisy, we register the k^{th} component $C_k(T_{0.40})$ with the k^{th} component $C_k(I_{0.40})$. For each registration simulation, denote by ψ the optimal transformation produced by the registration algorithm, and denote by ψ_X and ψ_Y the corresponding X - and Y -translation parameters of the optimal transformation ψ . We present the results of this multiscale registration in Table 7.

To estimate the transformation Ψ , for mutual information we average the parameters corresponding to registration of the first 7 scales. For mean squares, we average the results of the first 2 registrations, and for normalized correlation, we average the registration results from the first 4 registrations. In Table 8, we present the X - and Y -translation values Ψ_X and Ψ_Y .

Note that since the actual translation values are 13 in X and 17 in Y , our multiscale registration technique provides accurate results in the case in which both the fixed and moving images contain significant levels of noise.

REMARKS.

1. *For the sake of brevity, we presented only the multiscale registration results for registration of images that contain levels of noise greater than the level at*

TABLE 7. The results of registering $T_{0.40}$ with $I_{0.40}$ using Algorithm II. Here, we use $m = 12$ hierarchical steps to decompose the noisy image, so we perform $m = 12$ registration simulations. The actual translation values are 13 in X and 17 in Y .

Fixed Image	Moving Image	Mean Squares		Normalized Correlation		Mutual Information	
		ϕ_X	ϕ_Y	ϕ_X	ϕ_Y	ϕ_X	ϕ_Y
$I_{0.40}$	$T_{0.40}$	5.03	2.54	11.02	7.04	3.05	0.99
$C_0(I_{0.40})$	$C_0(T_{0.40})$	13.06	16.92	13.05	16.92	13.05	16.92
$C_1(I_{0.40})$	$C_1(T_{0.40})$	13.05	16.93	13.02	16.22	13.06	16.92
$C_2(I_{0.40})$	$C_2(T_{0.40})$	13.03	16.93	8.11	5.29	13.02	16.27
$C_3(I_{0.40})$	$C_3(T_{0.40})$	13.02	16.94	5.40	12.19	13.02	16.25
$C_4(I_{0.40})$	$C_4(T_{0.40})$	13.02	16.94	2.20	8.00	2.23	5.09
$C_5(I_{0.40})$	$C_5(T_{0.40})$	13.01	16.93	26.76	1.21	1.17	7.00
$C_6(I_{0.40})$	$C_6(T_{0.40})$	12.99	16.81	23.83	4.11	1.22	2.17
$C_7(I_{0.40})$	$C_7(T_{0.40})$	7.05	6.08	0.20	3.15	0.20	4.15
$C_8(I_{0.40})$	$C_8(T_{0.40})$	6.78	5.05	6.04	2.09	6.04	6.05
$C_9(I_{0.40})$	$C_9(T_{0.40})$	3.05	1.02	9.98	1.10	5.06	10.01
$C_{10}(I_{0.40})$	$C_{10}(T_{0.40})$	12.20	14.01	-1.97	0.99	-3.93	3.04
$C_{11}(I_{0.40})$	$C_{11}(T_{0.40})$	4.80	3.19	1.01	5.98	3.91	0.72

TABLE 8. The translation parameters Ψ_X and Ψ_Y for registration of $T_{0.40}$ with $I_{0.40}$ obtained by averaging the translation parameters of the coarse scale registrations. The actual translation values are 13 in X and 17 in Y .

	Mean Squares	Normalized Correlation	Mutual Information
Ψ_X	13.03	13.03	13.04
Ψ_Y	16.92	16.57	16.59

which ordinary registration methods fail. However, we also performed multiscale registration simulations for noise densities lower than those presented here, and in all cases, the multiscale technique was either as accurate as or more accurate than ordinary registration techniques.

2. The method of estimating the translation parameters in X and Y by averaging the parameters corresponding to the coarse scale registrations is based on determining the scales that should be taken into consideration. For the results presented in this paper, as well as for all of the other simulations that we studied, we found a drastic jump in the translation parameters such as that between the eighth and ninth scales in the mutual information column of Table 4. In such cases, the natural choice is to average the parameters corresponding to the coarse scales before the jump, and to exclude the remaining values; indeed, in all cases considered for this study, we have found that estimating the parameters in this way yields extremely accurate results. More generally, we expect that for most problems of this type, there should be a noticeable jump in the multiscale registration parameters, thus enabling a determination of the

coarse scales that should be averaged. This jump occurs because once a certain level of detail, and hence noise, appears in the scales, the registration process fails. More specifically, the optimization of the registration metric does not produce meaningful results if enough noise is present, and the registration parameters that result differ significantly from those that result from registration of images in which less noise is present.

8. Summary. While there are many existing medical image registration techniques, common approaches are shown to fail to give accurate results when one or more of the images to be registered contains high levels of noise. Further, if the noise level is significantly high, image registration can fail even when a denoising algorithm is applied to the noisy images before registration. We have presented an image registration technique based on the hierarchical multiscale decomposition of [20] of the images to be registered. The multiscale decomposition of an image results in a hierarchical representation that separates the coarse and fine scales of the image. Upon obtaining the decomposition of one or both of the images to be registered, we register the components of the moving image with the components of the fixed image. Since the coarse scale components of an image contain the essential features and shapes of the image, registration of the coarse scale components of the moving image with the coarse scale components of the fixed image provides an accurate estimate of the actual transformation that brings the moving image into spatial alignment with the fixed image. Using images in which the precise transformation that maps one to the other is known, we have shown that the multiscale approach is indeed accurate for levels of noise much higher than the noise levels at which ordinary optimal linear registration and denoising prior to ordinary registration methods fail; moreover, for all levels of noise, the multiscale technique either matches or outperforms ordinary registration techniques. Finally, we hope to extend these techniques to other image registration problems in which ordinary registration techniques are not successful due to degradation or other factors present in the images to be registered.

Acknowledgments. The work of D. Levy was supported in part by the National Science Foundation under Career Grant No. DMS-0133511. The work of L. Xing was supported in part by the Department of Defense under Grant No. PC040282 and the National Cancer Institute under Grant No. 5R01 CA98523-01.

REFERENCES

- [1] N.M. ALPERT, J.F. BRADSHAW, D. KENNEDY, AND J.A. CORREIA, *The principal axes transformation-A method for image registration*, Journal of Nuclear Medicine, vol. 31, no.10, pp. 1717-1722, 1990.
- [2] M. BRO-NIELSEN AND C. GRAMKOW, *Fast fluid registration of medical images*, Lecture Notes in Computer Science, 1131, pp. 267-276, Springer, Berlin and Heidelberg, 1996.
- [3] C. BROIT, *Optimal Registration of Deformed Images*, Ph.D. thesis, Computer and Information Science, University of Pennsylvania, 1981.
- [4] A. COLLIGNON, D. VADERMEULEN, P. SUETENS, AND G. MARCHAL, *3d multi-modality medical image registration based on information theory*, Computational Imaging and Vision, vol. 3, pp. 263-274, 1995.
- [5] R. DUDA AND P. HART, *Pattern Classification and Scene Analysis*, John Wiley and Sons, New York, 1973.
- [6] B. FISCHER AND J. MODERSITZKI, *Fast inversion of matrices arising in image processing*, Numerical Algorithms, vol. 22, pp. 1-11, 2001.
- [7] *Insight Segmentation and Registration Toolkit (ITK) Software Guide*, <http://www.itk.org>.

- [8] W.A. LIGHT, *Variational methods for interpolation, particularly by radial basis functions*, Numerical Analysis, pp. 94-106, 1995.
- [9] F. MAES, A. COLLIGNON, D. VANDERMEULEN, G. MARCHAL, AND P. SUETENS, *Multimodality image registration by maximization of mutual information*, IEEE Transactions on Medical Imaging, vol. 16, no. 2, pp. 187-198, 1997.
- [10] S. MALLAT, *A Wavelet Tour of Signal Processing*, Academic Press, Paris, 1998.
- [11] Y. MEYER, *Oscillating Patterns in Image Processing and Nonlinear Evolution Equations*, University Lecture Series 22, AMS, Providence, RI, 2002.
- [12] J. MODERSITZKI, *Numerical Methods for Image Registration*, Oxford, 2004.
- [13] P. PERONA AND J. MALIK, *Scale-space and edge detection using anisotropic diffusion*, IEEE Transactions on Pattern Analysis and Machine Intelligence, vol. 12, no. 7, pp. 629-639, 1990.
- [14] T. M. PETERS, B. L. K. DAVEY, P. MUNGER, R. M. COMEAU, A. EVANS, AND A. OLIVER, *Three-dimensional multi-modal image-guidance for neurosurgery*, IEEE Transactions on Medical Imaging, vol. 15, no. 2, pp. 121-128, 1996.
- [15] J. P. W. PLUIM, J. B. A. MAINTZ, AND M. A. VIERGEVER, *Image registration by maximization of combined mutual information and gradient information*, IEEE Transactions on Medical Imaging, vol 19, no. 8, pp. 809-814, 2000.
- [16] J. P. W. PLUIM, J. B. A. MAINTZ, AND M. A. VIERGEVER, *Mutual information based registration of medical images: A survey*, IEEE Transactions on Medical Imaging, vol. 22, no. 8 , pp. 986-1004, 2003.
- [17] K. ROHR, *Landmark-based Image Analysis*, Computational Imaging and Vision, Kluwer Academic, Dordrecht, 2001.
- [18] L. RUDIN AND V. CASELLES, *Image recovery via multiscale total variation*, Proceedings of the Second European Conference on Image Processing, Palma, Spain, 1995.
- [19] L. RUDIN, S. OSHER, AND E. FATEMI, *Nonlinear total variation based noise removal algorithms*, Physica D, vol. 60, pp. 259-268, 1992.
- [20] E. TADMOR, S. NEZZAR, AND L. VESE, *A multiscale image representation using hierarchical (BV, L^2) decompositions*, Multiscale Modeling and Simulations, vol. 2, no. 4, pp. 554-579, 2004.
- [21] P. THÉVENAZ AND M. UNSER, *Optimization of mutual information for multiresolution image registration*, IEEE Transactions on Image Processing, vol. 9, no. 12, pp. 2083-2099, 2000.
- [22] P. VIOLA, W. WELLS, H. ATSUMI, S. NAKAJIMA, AND R. KIKINIS, *Multi-modal volume registration by maximization of mutual information*, Medical Image Analysis, vol. 1, no. 1, pp. 35-51, 1995.
- [23] A. R. WEEKS, *Fundamentals of Electronic Image Processing*, SPIE Optical Engineering Press and IEEE Press, 1996.
- [24] J. WEST, ET AL., *Comparison and evaluation of retrospective intermodality brain image registration techniques*, Journal of Computer Assisted Tomography, vol. 21, no. 4, pp. 554-566, 1997.

Received on December 29, 2005. Accepted on January 26, 2006.

E-mail address: dpaquin@stanford.edu

E-mail address: dlevy@math.stanford.edu

E-mail address: edwards@stanford.edu

E-mail address: lei@reyes.stanford.edu



OVERVIEW OF IMAGE-GUIDED RADIATION THERAPY

LEI XING, PH.D., BRIAN THORNDYKE, PH.D., EDUARD SCHREIBMANN, PH.D.,
YONG YANG, PH.D., TIAN-FANG LI, PH.D., GWE-YA KIM, PH.D., GARY LUXTON, PH.D.,
and ALBERT KOONG, M.D.

Department of Radiation Oncology, Stanford University School of Medicine, Stanford, CA

(Accepted 21 December 2005)

Abstract—Radiation therapy has gone through a series of revolutions in the last few decades and it is now possible to produce highly conformal radiation dose distribution by using techniques such as intensity-modulated radiation therapy (IMRT). The improved dose conformity and steep dose gradients have necessitated enhanced patient localization and beam targeting techniques for radiotherapy treatments. Components affecting the reproducibility of target position during and between subsequent fractions of radiation therapy include the displacement of internal organs between fractions and internal organ motion within a fraction. Image-guided radiation therapy (IGRT) uses advanced imaging technology to better define the tumor target and is the key to reducing and ultimately eliminating the uncertainties. The purpose of this article is to summarize recent advancements in IGRT and discussed various practical issues related to the implementation of the new imaging techniques available to radiation oncology community. We introduce various new IGRT concepts and approaches, and hope to provide the reader with a comprehensive understanding of the emerging clinical IGRT technologies. Some important research topics will also be addressed. © 2006 American Association of Medical Dosimetrists.

Key Words: IGRT, Organ motion, Image guidance, Dose optimization, 4D imaging.

INTRODUCTION

Radiotherapy is an image-guided intervention, and imaging is involved in every key step of the process, ranging from patient staging, simulation, treatment planning, and radiation delivery, to patient follow-up. The evolution of radiation therapy has been strongly correlated with the development of imaging techniques. During the early days when Roentgen first discovered x-rays, 2-dimensional (2D) transmission images of the human body provided unprecedented imagery of bony landmarks, which allowed radiologists to deduce the location of internal organs. Using planar radiographs, radiologists planned cancer treatments by collimating rectangular fields that circumscribed the presumed tumor location. Additional blocks placed daily to match marks on the patient's skin, and later using low-temperature-melting dense alloys. The emergence of computed tomography (CT) in the 1970s revolutionized radiation therapy and allowed us to use image data to build a 3-dimensional (3D) patient model and design 3D conformal radiation treatment. In general, 3D conformal radiation therapy (3DCRT) is a method of irradiating a tumor target volume defined in a 3D anatomical image of the patient with a set of x-ray beams individually shaped to conform to the 2D beam's-eye-view (BEV) projection of the target. The reduction in normal tissue irradiation when moving

from 2D to 3D should theoretically improve the therapeutic ratio and allow the tumor target volume to be treated to a higher dose, thereby improving the probability of tumor control. Recent technical advances in planning and delivering intensity-modulated radiation therapy (IMRT) provide an unprecedented means for producing exquisitely shaped radiation doses that closely conform to the tumor dimensions while sparing sensitive structures.^{1–3} The development of 3DCRT and IMRT places more stringent requirements on the accuracy of beam targeting. In practice, large uncertainties exist in tumor volume delineation and in target localization due to intra- and inter-organ motions. The utility of modern radiation technologies, such as 3DCRT and IMRT, cannot be fully exploited without eliminating or significantly reducing these uncertainties. The need to improve targeting in radiation treatment has recently spurred a flood of research activities in image-guided radiation therapy (IGRT).

While all radiation therapy procedures are image guided *per se*, traditionally, imaging technology has primarily been used in producing 3D scans of the patient's anatomy to identify the location of the tumor prior to treatment. The verification of a treatment plan is typically done at the level of beam portals relative to the patient's bony anatomy before patient treatment. In current literature, the term of IGRT or IG-IMRT is employed loosely to refer to newly emerging radiation planning, patient setup, and delivery procedures that integrate cutting-edge image-based tumor definition methods, pa-

Reprint requests to: L. Xing, Ph.D., Department of Radiation Oncology, Stanford University School of Medicine, 875 Blake Wilbur Drive, Stanford, CA 94305-5847. E-mail: lei@reyes.stanford.edu

tient positioning devices, and/or radiation delivery guiding tools. These techniques combine new imaging tools, which interface with the radiation delivery system through hardware or software, and state-of-the-art 3DCRT or IMRT, and allow physicians to optimize the accuracy and precision of the radiotherapy by adjusting the radiation beam based on the true position of the target tumor and critical organs. With IGRT, it is also possible to take tumor motion into account during radiation therapy planning and treatment. Because IGRT improves precision, it raises the possibility of shortening the duration of radiation therapy by reducing the number of treatment sessions for some forms of cancer.

The purpose of this article is to highlight the recent developments of various available imaging techniques and present an overview of IGRT. Stanford experience on various aspects of clinical IGRT will also be presented. We hope that readers will gain an overall picture of IGRT and find it easier to navigate themselves through the subsequent articles in this issue, which focus on providing technical details and/or specific clinical applications of the available IGRT tools.

ISSUES IN IGRT

In current 3DCRT or IMRT, uncertainties exist in many circumstances, such as tumor target definition, patient immobilization, and patient breathing motion, which make it difficult to administer a high radiation dose to the planned location. The exact locations of the boundaries of the tumor target and the adjacent sensitive structures are often not known precisely, and a population- and disease site-based safety margin is used routinely to cope with a problem that is otherwise insoluble. An important task of IGRT is to eliminate or significantly reduce the margins involved in defining the clinical and planning target volume (CTV and PTV, respectively).

Many IGRT solutions have been proposed to resolve the problem of target definition and beam targeting. Briefly, IGRT developments are focused in four major areas: (1) biological imaging tools for better definition of tumor volume; (2) time-resolved (4D) imaging techniques for modeling the intra-fraction organ motion; (3) on-board imaging system or imaging devices registered to the treatment machines for inter-fraction patient localization; and (4) new radiation treatment planning and delivery schemes incorporating the information derived from the new imaging techniques. These are discussed in more detail in the following.

TUMOR TARGET VOLUME DEFINITION

CT, MRI, and ultrasound (US) imaging techniques

To be able to "see" the extent of disease more clearly and define the tumor target volume relative to the patient's anatomy have been among the most important issues in radiation oncology. CT has played a pivotal role

in the process. Many radiation oncology departments have acquired dedicated CT scanners. A typical patient's 3D CT data set has more than 100 axial slices, each of which contains 512×512 pixels. With 16 bits per pixel, a CT data set can easily run over 50 megabytes. CT has many advantages, including high spatial integrity, high spatial resolution, excellent bony structure depiction, and the ability to provide relative electron density information used for radiation dose calculation. The recent development of ultra-fast multi-slice CT has opened a new dimension to CT technology and allows time-resolved (4D) CT imaging of patient's cardiac and breathing cycles. Using array detectors, multisector CT scanners can acquire multiple slices or sections simultaneously and thereby greatly increase the speed of CT image acquisition. Currently, all manufacturers are moving toward 8-, 16- and even higher slice CT technology. Radiation oncology application of 4D CT will be discussed later.

MRI provides superior soft tissue discrimination, especially for central nervous system (CNS) structures and within the abdomen and pelvis, and has been widely used in the diagnosis and tumor delineation. MRI is also utilized for virtual simulation of radiation treatment for some specific disease sites. Physically, MRI involves the determination of the bulk magnetization of nuclei within a given voxel through use of radio-frequency (RF) radiation and magnetic fields. In a clinical setting, MRI is typically employed together with CT images with the help of image fusion software to delineate the extent of the malignancy. As with other imaging techniques, MR technology has gone through a series of revolutions in the past 3 decades. MRI technology is moving toward higher field strengths to further improve the quality of MR images, as evidenced by the installations of 3T scanners in many institutions (9.4 T MRI scanners have been installed in a few institutions). Fast-cine MRI is also becoming increasingly available and may offer physicians an alternative for imaging the temporal process of patient breathing or even heart beating. Figure 1 shows an example of MRI images acquired at 2 different phases for a liver cancer patient. In addition, the development of some specialized MRI scans has also attracted much attention. These include diffusion and perfusion MRI, dynamic contrast MRI, MR angiography, MR spectroscopic imaging (MRSI), and functional MRI (fMRI). The recent development of diffusion tensor imaging (DTI), for instance, enables diffusion to be measured in multiple directions and the fractional anisotropy in each direction to be calculated for each voxel. fMRI measures signal changes in the brain that are due to changing neural activity. These techniques enable researchers to make axonal and functional maps to examine the structural connectivity of different regions in the brain and may allow better definition of brain tumors and better sparing of sensitive regions.⁴

Ultrasound (US) is another useful imaging modality for radiation therapy. US utilizes high-frequency (1~10

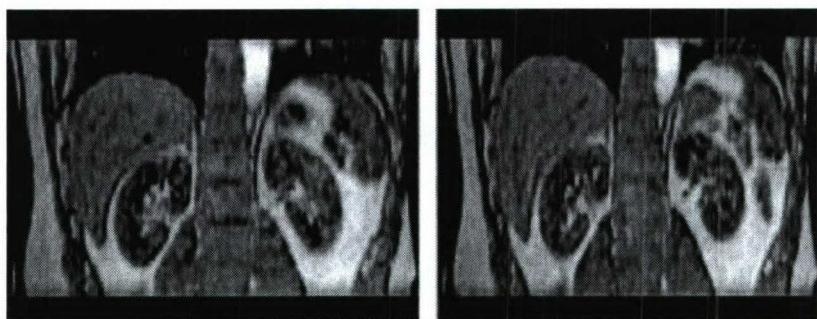


Fig. 1. Cine MR images at inhale and exhale phases for a liver cancer patient.

MHz) sound waves to generate anatomical images that have high spatial resolution and tissue characterization discrimination power through image texture analysis. In radiation therapy, it has been particularly useful in prostate imaging. Transrectal US permits an examination/localization of the prostate gland^{5,6} and is the imaging modality of choice in guiding the prostate seed implant procedure.

Biological imaging

Regardless of the course of therapy, current standard imaging modalities such as CT and MRI do not always provide an accurate picture of the tumor extent, especially in the zone of infiltration that may be the limiting factor in an attempt of a radical treatment approach. This has been shown to be the case for gliomas before surgical intervention. It is also true when attempting to determine the volume of residual tumor for additional therapy owing to problems in differentiating post-therapy changes from residual tumor. Indeed, the above-mentioned imaging modalities are anatomic in nature, *i.e.*, they provide snapshot of a patient's anatomy without biological information of various organs or structures. Biological imaging, defined as the *in vivo* characterization and measurement of biological processes at the cellular and molecular level, is an emerging multidisciplinary field resulting from the developments of molecular biology and diagnostic imaging and shows significant promise to revolutionize cancer detection, staging/re-staging, treatment decision-making, and assessment of therapeutic response. MRSI and positron emission tomography (PET) are 2 valuable modalities for radiation therapy planning. ¹H MRSI combines the advantages of obtaining biochemical data by water-suppressed ¹H MR spectroscopy with the spatial localization of that data. MR spectroscopy is useful in characterization of brain and prostate tumors. In the brain, for example, malignant tumors have an increased rate of membrane turnover (increased level of choline) and a decreased concentration of neurons. Furthermore, spectroscopy allows for the noninvasive monitoring of the response of residual tumor to therapy and for differentiating tumor recurrence

from tissue necrosis. Recently, Pirzkall *et al.*⁷ have applied multi-voxel MRSI to assess the impact of MRSI on the target volumes used for radiation therapy treatment planning for high-grade gliomas. It was found that, although T₂-weighted MRI estimated the region at risk of microscopic disease as being as much as 50% greater than by MRSI, metabolically active tumor tissue still extended outside the T₂ region in 88% of patients by as much as 28 mm. In addition, T₁-weighted MRI suggested a lesser volume and different location of active disease compared to MRSI. The discordance of high-grade-glioma target volumes resulting from MRI was also observed in other functional imaging modalities such as (PET) and single-photon emission computed tomography (SPECT).

While there is a growing body of evidence now indicating that *in vivo* MRSI provides unique information on metabolism that will ultimately affect clinical diagnosis, choice and monitoring of therapies, and treatment planning, in reality, MRSI has mainly remained a research tool confined to a small number of academic institutions.^{8–12} PET, on the other hand, is more widely used and has been harnessed into the planning process in many clinics. In general, PET has lower image resolutions than CT images and, with commonly used fluorine-18-labeled deoxyglucose (FDG) tracer, contains no anatomic information about normal structures. Information derived from PET needs to be fused with the corresponding CT images for treatment planning. The fusion of PET and CT images are simplified with the use of the hybrid PET/CT scanner.^{13,14} Figure 2 shows the data flow of a typical PET/CT scanner.

Hybrid PET/CT systems have several positive features that are absent in stand-alone PET and CT units. PET/CT is a hardware-based image-fusion technology that virtually eliminates the uncertainty and inconvenience of currently available software fusion of separate PET and CT images, which are often acquired with patients in different positions. It should be emphasized that the PET/CT unit is not simply a PET and CT combination—not from the perspective of system design, nor the practical utility. Other than the fact that one does

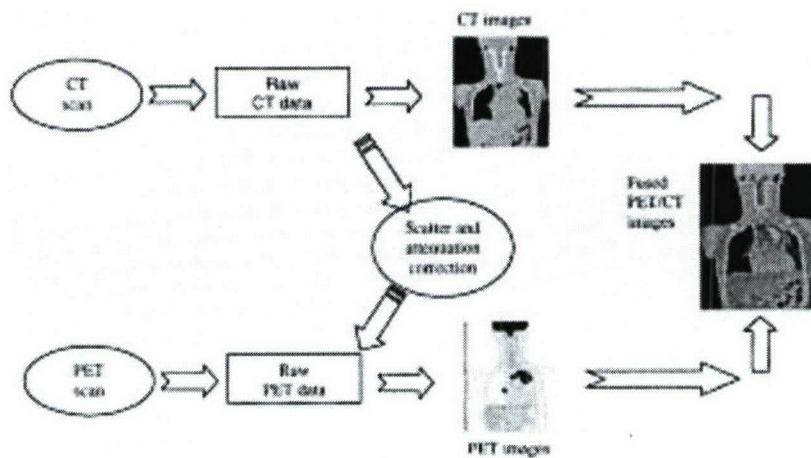


Fig. 2. Schematic drawing of the data flow in a hybrid PET/CT.

not have to go through the cumbersome and time consuming software fusion process, it has the advantages of simultaneous availability of the fused images, convenience to the patient and the physician, increased physician confidence in interpreting the image findings, and ~30% of reduction in PET scanning time due to the use of CT data for PET attenuation correction.

Integration of biological imaging techniques and multimodality image fusion

FDG-PET provides a means to study metabolic activity of tumors *in vivo*. Initial studies incorporating FDG-PET into treatment planning have been reported.^{15–17} Bradley *et al.*¹⁷ have carried out a prospective study to determine the impact of functional imaging with FDG-PET on target volumes among non-small cell lung cancer (NSCLC) patients being considered for definitive radiation therapy. They found that radiation targeting with fused FDG-PET and CT images resulted in alterations in radiation therapy planning in over 50% of patients by comparison with CT targeting. The changes included the alterations in the AJCC TNM stage (31% of the patients studied) and modification of target volume (58% of the patients studied). In a separate study, Mac-Manus *et al.*¹⁶ reported that 30% of patients with locally-advanced NSCLC became ineligible for curative radiotherapy because of detection of either distant metastatic disease or intrathoracic disease too extensive for radical radiation. Recently, Howard *et al.*¹⁸ have studied the value of FDG-PET/CT for esophagus cancer and reported similar findings.

Emerging PET tracers for oncologic imaging

While FDG-PET has been shown to be effective for a number of malignancies, imaging of many other neoplasms, such as breast cancer and prostate cancer, with FDG has shown less success.^{19,20} Many pitfalls have previously been described with FDG-PET imaging. The

FDG tracer can be nonspecifically taken up by several benign conditions such as inflammatory disease, pneumonia, brown fat, muscle, bowel uptake, and granulomatous disease. Also, slow-growing indolent tumors may exhibit only a mild increase in glucose metabolism and therefore be missed by FDG-PET.^{21–23} Thus, FDG-PET is only minimally useful for the evaluation of indolent tumors such as organ-confined prostate cancer. The recent development of fluorothymidine (FLT)^{24–26} provided a new opportunity to improve the sensitivity and specificity of PET imaging of cancer. Because there is upregulation of thymidine transport into malignant cells due to accelerated deoxyribonucleic acid synthesis, either ¹¹C or ¹⁸F-labeled thymidine radiotracers can be used to determine cellular proliferation. Several studies have shown that the accumulation of FLT correlates better with proliferation in comparison with the commonly used FDG tracer.^{25,26} Recently, Smyczek-Gargya *et al.*²⁷ have reported FLT-PET imaging experiments involving 12 patients with 14 primary breast cancer lesions (T2–T4). Thirteen of the 14 primary tumors demonstrated focally increased FLT uptake. The study showed that FLT-PET is suitable for the diagnosis of primary breast cancer and locoregional metastases and the high image contrast of the technique may facilitate the detection of small foci.

Agents, such as antisense molecules, aptamers, antibodies, and antibody fragments, can be aimed at molecular targets for biological imaging. Tumor receptors and certain cellular physiologic activities, including metabolism, hypoxia, proliferation, apoptosis, angiogenesis, and infection, provide such targets. In addition to FLT, there are several other new nuclide imaging tracers under clinical or laboratory investigations,^{21,28–35} which include, to name a few, ¹¹C-Acetate,^{36–38} ¹⁸F-choline,^{39,40} ¹¹C-choline,^{41,42} ⁶⁴Cu-DOTA-Bombesin⁴³, ¹⁸F-FMISO,^{44,45} ¹⁸F-FAZA,⁴⁶ ⁶⁴Cu-ATSM.⁴⁷ For example, carcinogenesis is often characterized by enhanced cell proliferation

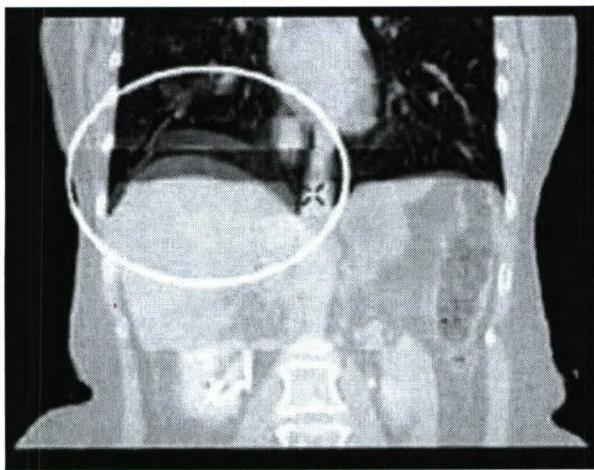


Fig. 3. Example of 4D CT where respiratory cycle irregularities have produced significant interbed mismatches near the base of the lung.

and transformation, and elevated levels of choline and choline kinase activity in certain neoplastic diseases have motivated the development of positron-labeled choline analogs for noninvasive detection of cancer using PET.⁴¹ Choline acts as a precursor for the biosynthesis of phospholipids, *e.g.*, phosphatidylcholine, the major components of cell membrane. Several preliminary studies have demonstrated the potential of the new tracer for prostate cancer and many other cancers.^{37,41,48}

Biologically conformal radiation therapy (BCRT)

The current 3DCRT or IMRT inverse planning is typically aimed at producing a homogeneous target dose under the assumption of uniform biology within the target volume. In reality, it is well known that the spatial biology distribution (*e.g.*, clonogen density, radiosensitivity, tumor proliferation rate, functional importance) in most tumors and sensitive structures is heterogeneous. Recent progress in biological imaging is making the mapping of this distribution increasingly possible. This new development opens a new avenue of research, coined BCRT.^{49–53} The goal of BCRT is to take the inhomogeneous biological information derived from biological imaging into account and to produce customized nonuniform dose distributions on a patient specific basis. The simultaneous integrated boost (SIB) to elective volumes recently appearing in the literature represents a simple example of BCRT.

To establish BCRT, 3 major aspects must be addressed: (1) determination of the distribution of biological properties of the tumor and critical structures; (2) prescription of the desired dose distribution for inverse planning; and (3) inverse planning to generate most faithfully the prescribed nonuniform dose distribution. While the development of molecular imaging techniques is critically important in mapping biology distributions,

the successful integration of this information into IMRT planning through steps (2) and (3) is also indispensable to fully exploit the obtained biology information to improve patient care. With the optimistic assumption that spatial biology distributions within a patient can be reliably determined using biological imaging in the future, Yang and Xing⁵³ have established a theoretical framework to quantitatively incorporate the spatial biology data into IMRT inverse planning. To implement this method, they first derived a general formula for determining the desired dose to each tumor voxel for a known biology distribution of the tumor based on a linear-quadratic (LQ) model. By maximizing the TCP under the constraint of constant integral target dose, they obtained

$$D_0^T(i) = \frac{\alpha'_ref}{\alpha'_i} D_{ref} - \frac{1}{\alpha'_i} (\gamma_{ref} - \gamma_i) \Delta T - \frac{1}{\alpha'_i} \ln \left(\frac{\alpha'_ref \rho_{ref}}{\alpha'_i \rho_i} \right), \quad (1)$$

where $D_0^T(i)$ is the desirable prescription dose at the voxel i with the tumor cell density, radiosensitivity, and proliferation rate given by $(\rho_i, \alpha_i, \gamma_i)$, and D_{ref} is the reference dose for the voxel with reference radiobiological parameters $(\rho_{ref}, \alpha_{ref}, \gamma_{ref})$. For a given disease site, the radiation dose used in current clinical practice with “intent to cure” can be used as a good starting point in selecting the value of D_{ref} . The relation is quite general and can be used as prescription dose to guide an arbitrary inverse planning objective function aimed at producing a customized dose distribution in accordance with the spatial biology information.

INTRAFRACTION ORGAN MOTION: MANAGING THE RESPIRATORY MOTION

Components affecting the reproducibility of target position during and between subsequent fractions of radiation therapy include the displacement of internal organs between fractions and internal organ motion within a fraction. Depending on the disease site, these components contribute differently to the margins that are to be added around the CTV to ensure adequate coverage. In the thorax and abdomen, intra-fraction internal anatomy motion due to respiration is a principal cause for large safety margins. Motion can distort target volumes and result in positioning errors as different parts of the tumor move in and out of the image window with the patient’s breathing cycle. Several studies, conducted to examine the extent of diaphragm excursion due to normal respiration, reported the range of motion from ~0.5 to 4.0 cm in the superioinferior direction. As a consequence of a significant margin added around the CTV, a large amount of normal tissue surrounding the CTV is irradiated. Accounting for such motion during treatment has the potential to reduce margins, leading to reduced radiation toxicity and risk of treatment-induced complications, and yielding room for dose escalation.

A complete solution compensating for respiratory motion should ideally start at the simulation stage. There have been several studies to characterize the amplitude, phase and periodicity of organ motion^{54–56} using fluoroscopic x-rays, ultrasound,^{57,58} and magnetic or RF markers.^{59,60} The development and deployment of spiral and multi-detector CT scanners have made practical the acquisition of time-resolved or 4D CT images. The reconstructed images acquired with patients in treatment positions provide 4D models upon which geometric as well as dosimetric computations can be performed. 4D PET is also becoming clinically available.^{61–63} Treatment-wise, respiratory gating technology and tumor tracking techniques to synchronize delivery of radiation with the patient's own respiratory cycle are under intensive investigations.

4D CT imaging

A 4D CT can be either prospective or retrospective. In the former case, the scanner collects images at only one of the breathing phases of the patient instead of scanning continuously. The retrospective 4D CT scan results in multiple image sets, corresponding to different breathing phases of the patient, and consists of 3 relatively orthogonal processes^{64–68}: recording of respiratory signal(s), acquisition of time-dependent CT projection data, and construction of a 4D image from these data. The first objective can be achieved by tracking a surrogate of respiration-related organ and tumor motion, such as tidal volume measured with a spirometer,^{66,69} chest expansion monitored by a pneumatic bellows,⁷⁰ or a reflecting external marker placed on the abdomen and tracked with a camera.⁶⁴ Time-dependent CT data can be acquired by oversampling in either helical or cine mode, and constructing several CT slices over the full respiratory cycle at each axial location.^{67,71} Finally, the respiratory signal and CT data must be combined into a 4D series, providing a CT volume at several points throughout the respiratory cycle. In this section, we will focus primarily on the implementation of 4D CT provided by the Varian Real-time Position Management (RPM) camera/software and the GE Discovery ST multislice PET/CT scanner.

4D CT patient setup proceeds along the same lines as a standard 3D CT exam. The patient is immobilized on the scanner bed, and aligned using room and scanner lasers. Sagittal and coronal scout images are used to verify patient positioning, and the setup is adjusted as necessary. At this stage of the setup, the 4D procedure begins to diverge from the 3D exam.

The RPM system consists of an infrared source, CCD camera, and a reflecting block. The block is attached to the patient's abdomen, typically just inferior to the xiphoid process, and the anteroposterior motion of the block is captured by the camera. This motion is analyzed in real-time by Varian software on a computer connected to the RPM camera. The breathing pattern is

recorded for the duration of the scan, and is referred to as the "respiratory trace." Once the scan has finished, the software retrospectively computes the phase at each point of the respiratory trace by determining the location of the peaks at end-inspiration, and assigning percentages to interpeak points based on a linear interpolation of the peak-to-peak distance. For example, under this scheme, end-inspiration occurs at 0%, while end-expiration typically appears near 50–60%. The peak-to-peak distance can vary between respiratory cycles, as can the position of end-expiration with respect to end-inspiration.

Irregularities in a patient's respiratory pattern can often be reduced by encouraging the patient to breathe calmly and consistently, and then relying on the patient's compliance during the scan. If this free-breathing approach is insufficient, the RPM software can provide audio coaching in the form of a "breathe in, breathe out" recording, which is manually or automatically timed to the patient's natural rhythm. Some groups have used video feedback either alone or concurrently with audio instructions.⁷² While audio and video coaching can help by stabilizing the respiratory period, amplitude and baseline, they can complicate matters for patients with compromised respiratory function, who find it difficult or impossible to maintain a regular rhythm. Another solution is active breath control (ABC)^{73–75} which uses modified ventilator equipment to control the airflow, albeit at the (possibly significant) expense of patient comfort.

Once a sufficiently regular breathing pattern has been established, the CT data is acquired in "cine" mode. This is a step-and-shoot technique, whereby the gantry rotates several times at each bed position to acquire data over the full respiratory cycle. The raw data is partitioned into bins corresponding to a user-selected time interval (typically less than 1/10th the average cycle), and CT slices are automatically reconstructed from these bins. Because several respiratory points are sampled at each bed position, a 4D CT scan can take several times as long as a corresponding 3D CT, resulting in typically 1500–3000 CT slices for a 20–40-cm axial FOV.

The respiratory and scan data are combined at a separate computer, the Advantage Workstation (AW) (GE Medical Systems), which uses the respiratory trace to sort the oversampled CT slices according to their phase. The AW does perform the phase calculations, but rather relies on the phase stamp computed by the RPM during the creation of the respiratory trace file. Missing phases for any couch position are replaced with their nearest neighbor, providing a sorted image without any phase gaps. The user can navigate through the data in each axial direction, similar to standard viewing software, but can also scroll through the respiratory phases from end-inspiration to end-expiration. Individual phases can be subsequently extracted, or combined into averaged or minimum/maximum intensity projections, and exported to planning software in the form of standard

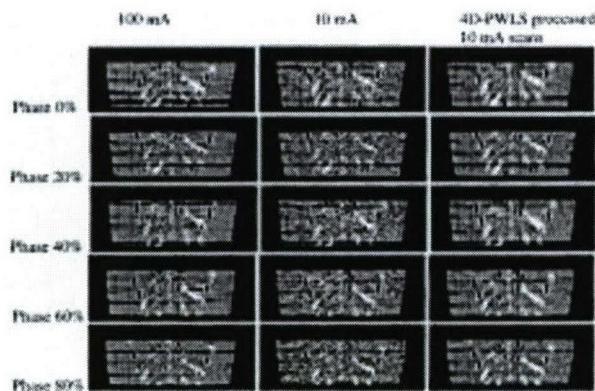


Fig. 4. Motion phantom study for the 4D-PWLS method with the thorax phantom. The left and middle columns are the original phases obtained from the GE Advantage Workstation, for 100 mA and 10 mA, respectively; the right column shows the 10-mA phases after 4D-PWLS processing. The red rectangles represent the selected ROI for calculation of SNRs, each of which contains $5 \times 5 \times 5$ voxels. PWLS smoothed 10-mA scan resulted in more than 2-fold increase in the SNR for every phase of the periodically moving phantom. Similar results were obtained in a patient 4D CT study.

DICOM series. These exported image series form the basis of 4D treatment planning.

Unresolved issues in 4D CT

The AW sorts the data by phase rather than amplitude. If the breathing were perfectly regular from cycle to cycle, then phase- and amplitude-based sorting would give very similar results. The problem arises when there is variation in amplitude, period, or baseline, or when the onset of end-expiration does not occur at the same point each cycle. When these inconsistencies arise, the sorted CT images may contain mismatch artifacts at the interface between bed positions (see Fig. 3). Recent studies have investigated amplitude-based binning as an alternative to the phase-based approach, and it appears that amplitude sorting can improve image quality in many cases.^{76–78} Other researchers have matched adjacent CT slices without using a respiratory trace, by maximizing the continuity of CT units integrated over regions of interest.⁷¹ Yet another promising approach involves interpolating the CT data continuously between end-cycle peaks using deformable models.⁷⁹

A second issue arises in the correlation between external fiducial movement and tumor/organ motion. Amplitude ratios between fiducial and tumor displacement may vary from cycle to cycle, and thoracic and abdominal points may involve relative phase shifts.^{54,80} These shifts may be especially crucial for tumors near the lung, where hysteresis is prevalent. Finally, larger organs such as the liver can experience substantial deformation during inspiration and expiration, which may not be adequately captured by rigid-body interpolation between points in the respiratory cycle.

Finally, even if the 4D CT images have been acquired without problem, there remains the issue of reproducibility at treatment.⁸¹ If treatment planning and delivery are based on 4D CT, there is an implicit assumption that anatomic motion during treatment will match the tumor and organ motion observed during setup. This assumption can be checked to some degree through frequent gated or breath-hold portal imaging.⁸² On the other hand, it is reasonable to assume the patient will relax over time, so that their breathing becomes shallower or changes tempo. Indeed, studies have demonstrated that some patients exhibit systematic respiratory changes over a several-week course of radiation therapy, even with visual and audio coaching.⁸³ These issues strike at the heart of IGRT, and provide a fertile ground for research.

4D CT usually delivers more radiation dose than the standard 3D CT, because multiple scans at each couch position are required to provide the temporal information. We have developed a method to perform 4D CT scans at relatively low current, hence reducing the radiation exposure of patients.⁶⁸ To deal with the increased statistical noise caused by the low current, we proposed a novel 4D penalized weighted least square (4D-PWLS) smoothing method, which can incorporate both spatial and phase information. The 4D images at different phases are registered to the same phase via a deformable model, whereby a regularization term combining temporal and spatial neighbors is designed. The proposed method was tested with phantom experiment (see Fig. 4), for an example) and patient study, and superior noise suppression and resolution preservation were observed.

4D PET and related issues

4D PET poses a problem distinct from 4D CT, in that signal is inherently limited by the tolerable patient dose. The result is that any PET scan requires a significant amount of time per bed position (usually a few minutes) to acquire sufficient data to produce a good image. This limitation makes it difficult to partition PET data with the same time resolution possible in 4D CT, but nonetheless, acquisition methods are clinically available to obtain PET images at end-inspiration or end-expiration. The most common solution is to gate the PET scan at the desired respiratory end-point, and reconstruct a single bin of gated data.^{84–86}

Patient setup proceeds in the same manner as an ungated PET scan, and a CT image is acquired for attenuation correction just prior to the PET. At this point, the RPM system monitors patient breathing by tracking the reflecting block, and the acquisition trigger is set by the user to occur at some given point (say, end-inspiration) in the cycle. Each time the RPM camera determines that the reflecting block (and, by extension, the patient's respiration) reaches this point in the respiratory cycle, a trigger is sent to the scanner, and data accumulation is initiated. Gated PET differs fundamen-

tally from the 4D CT protocol, by elevating the RPM system to this active role in data acquisition.

In gated mode, the user is able to select both the width of the acquisition window and the number of sequential bins to record each respiratory cycle. The bin width directly affects image quality, because the signal-to-noise ratio within an image asymptotically approaches the square root of the signal level.⁸⁷ Multiple bin acquisition allows the capture of the full respiratory cycle in several bins, offering the possibility of retrospectively sorting into 2 or more respiratory phases. Each time the RPM trigger is received, data is directed to the initial bin, and then to the remaining bins sequentially until the next trigger. This process continues for the duration of the scan. Ideally, the scan duration would be chosen such that the first bin (the respiratory point of most interest) would accumulate as many data points as a comparable ungated scan (*i.e.*, divide the bin width by the duty cycle). In reality, because this would lengthen the typical PET scan by a factor of 4 or 5, practical clinical considerations may require the gated scan to be shortened, with corresponding image degradation.

Once the scan has finished, it is possible to associate each bin (beyond the first bin) with a corresponding point in the respiratory cycle. Because the respiratory trace is recorded by the RPM, it is a relatively simple matter to analyze the respiratory motion offline and make this correspondence. It is also possible to retrospectively combine multiple bins into a single bin, merging all the data to create an effectively ungated scan. However, these methods are not yet available from the vendor as a clinical tool, and must be performed by the user in the context of research efforts. Once the desired bin has been selected, its data can be reconstructed using the vendor-supplied filtered backprojection or OS-EM algorithms. The image results can subsequently be exported to treatment planning systems for review, similar to ungated PET series.

A salient point in the PET reconstruction process is the specification of the attenuation correction map. The current clinical design uses the CT scan acquired just prior to the PET specifically for this purpose. This attenuation correction CT can be an acquired during either free breathing or breath-hold. Some research has indicated that PET reconstructions can be quite sensitive to distortions in the attenuation correction map,^{88–90} and investigations are ongoing into the use of 4D CT or other models to accurately account for attenuation.^{61,91} On the Varian/GE system, this requires selecting the appropriate images from the 4D CT on the AW, sending these series back to the scanner, generating the attenuation correction maps for each 4D PET bin, and then reconstructing each bin separately. Once again, this is a research solution, and not yet available from the vendor for clinical use.

Combining 4D PET with 4D CT and enhancement of the performance of 4D PET with post-acquisition data processing

Once the 4D PET has been acquired (either a single phase, or perhaps several), it is possible to create a 4D PET/CT.⁶¹ This involves manually selecting the PET and CT images with corresponding respiratory phases (or amplitudes), and fusing them on viewing/planning software. We have recently developed a 4D-4D image registration algorithm, which allows us to automate the process. If the CT and PET scans are acquired with the same patient position on the same exam, then the process is a particularly simple hardware-based registration. On the Eclipse treatment planning system, for example, 2 images (not just PET/CT, but other modalities as well) can be automatically fused if they share the same DICOM coordinates. If the DICOM coordinates are not identical, the registration is more difficult, requiring manual or automated shifts and rotations to match anatomical landmarks or fiducials. Fusion may be additionally complicated by organ deformation^{92,93} (see Rigid and Deformable Image Registration Section below). At the present time, PET/CT hardware fusion for ungated scans is well established and readily available within the clinical setting.^{13,14} 4D PET/CT registration, however, remains primarily within the research domain.

The major issue in 4D PET is the lack of statistics. Because the collected photons are divided into several frames, the quality of each reconstructed frame is decreased with increasing number of frames. The increased noise in each frame heavily degrades the quantitative accuracy of the PET imaging. We have recently developed 2 corrective methods to enhance the performance of 4D PET. The first method, coined “retrospective” stacking (RS),^{62,63,94} combines retrospective amplitude-based binning of data acquired in small time intervals, with rigid or deformable image registration methods. Unlike gating techniques, RS uses data along the entire respiratory cycle, thereby minimizing the need for lengthened scans while providing a 4-dimensional view of the region of interest.^{62,63} In the second approach,⁹⁴ we reconstruct each frame with all acquired 4D data by incorporating an organ motion model derived from 4D-CT images by modifying the well-known maximum-likelihood expectation-maximization (ML-EM) algorithm. During the processes of forward- and backward-projection in the ML-EM iterations, all projection data acquired at different phases are combined together to update the emission map with the aid of the deformable model, the statistics are therefore greatly improved. Both phantom and patient studies have indicated promising potential of the 2 methods.

Radiation treatment planning based on 4D information

How to maximally utilize the time-resolved image information derived from 4D CT or PET/CT represents one of the challenges in IGRT. In reality, the information

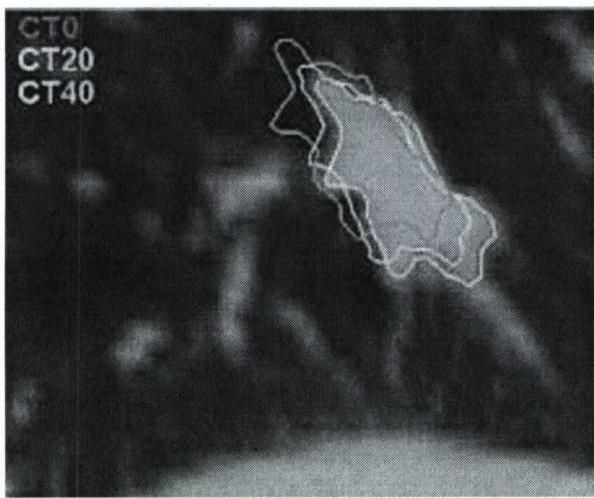


Fig. 5. Tumor contours for 3 breathing phases. The contours labeled as CT20 and CT 40 were produced by applying the deformation field on the tumor contours delineated in CT0. The "trajectories" of the tumor boundary pictorially show the extent of tumor movement and allow us to specify patient specific margin in accounting for the intra-fraction organ motion.

can be integrated into radiation treatment planning and delivery at different levels. At the lowest level, the 4D images can be employed to determine the extent of tumor movement on a patient specific basis and the information can then be used to design the CTV margin and the radiation portals to accommodate the motion. Figure 5 shows an example of lung patient, in which tumor boundaries at 3 distinct respiratory phases are plotted. We have referred to this type of treatment as "3.5-dimensional" radiation therapy. The 4D information can also be used for guiding breath-hold or gated radiation therapy. There is also strong interest in using the 4D data to establish a 4D patient model and then to carry out a 4D radiation therapy plan. These are the subjects of the following 2 sub-sections.

Breath-hold and respiratory gating

Various methods have been worked out to counteract respiratory motion artifacts in radiotherapy imaging. Among them are breath-hold, respiration gating, and 4D or tumor-tracking techniques.^{55,56,73,75,95} Breath-hold techniques either actively or passively suspend the patient's respiration and treat the patient during this interval. Deep inspiration breath-hold, active breathing control (ABC) (which forces shallow breathing and thereby "freezes" the tumor motion for a small part of the treatment time⁷³), and self-held breath-hold are suitable for different types of therapy targeting different cancers. Different types of equipment, such as stereotactic frames, fiducial tracking systems, timers, respirometer, RPM, or interlocks, may be needed depending on the method of breath-hold.

Respiration-gating methods involve tracking the pa-

tient's natural breathing cycle and periodically turning the beam on when the patient's respiration signal is in a certain phase of the breathing cycle (generally end-inhale or end-exhale). The patient's respiration is continuously monitored and the beam switches off as the tumor moves out of the target range. Gated radiation therapy can offset some of the motion but requires specific patient participation and active compliance. In gated treatment, it is required that the CT images used for treatment planning faithfully represent the actual treatment situation. While gated CT acquisition at the treatment respiratory phase is possible, our gating protocol proceeds by picking up the CT data at an appropriate phase from the patient's 4D CT acquired using the method described above. The gating interval is typically centered at end-expiration because of the increased reproducibility at this point, and spans 20–30% of the breathing period to provide a reasonable duty cycle. Treatment plans are optimized for this phase range by planning on an averaged composite of the scans within the interval, and using maximum- and minimum-intensity pixel views to incorporate intra-gate margins. The averaged, maximum-intensity and minimum-intensity composites for a lung patient are displayed in Fig. 6.

Tumor tracking

Similar to the establishment of a 3D geometric modeling based on traditional CT data, the availability of 4D imaging information makes it possible to build a patient specific 4D model. Figure 7 shows the the 4D model for a lung patient.⁷⁹ In obtaining the models, a BSpline deformable registration technique (see Rigid and Deformable Image Registration section below) was used to register different phases of the 4D CT. Ideally, organ motion represented by the 4D model can be incorporated into the radiation treatment plan optimization to overcome the adverse effect of respiratory motion on IMRT delivery.⁹⁶ A few groups^{95–100} have explored the feasibility of MLC-based tumor tracking. However, the interplay between different phases has been ignored during the plan optimization in most of these studies. Webb has presented a technique to model the dosimetric effect of elastic tissue movement when modulated beams are delivered.¹⁰¹ In general, the quadratic inverse planning objective function becomes

$$F = \sum_k \cdot \sum_i w_k \left[d_p^k - \sum_t d_i^k(\bar{r}, t) \right]^2 \quad (2)$$

where d_p^k is prescribed dose for k th structure, w_k is the importance factor, and $d_i^k(\bar{r}, t)$ is the calculated dose in voxel i at time t , and the summation over t represents the integral dose to i th voxel. For 4D planning, it is necessary to know the path of each material coordinate during the treatment, which involves registering the voxels in different respiratory phases. This can be achieved by using a deformable registration algorithm. The optimization of the above objective function or alike^{95,102–106}

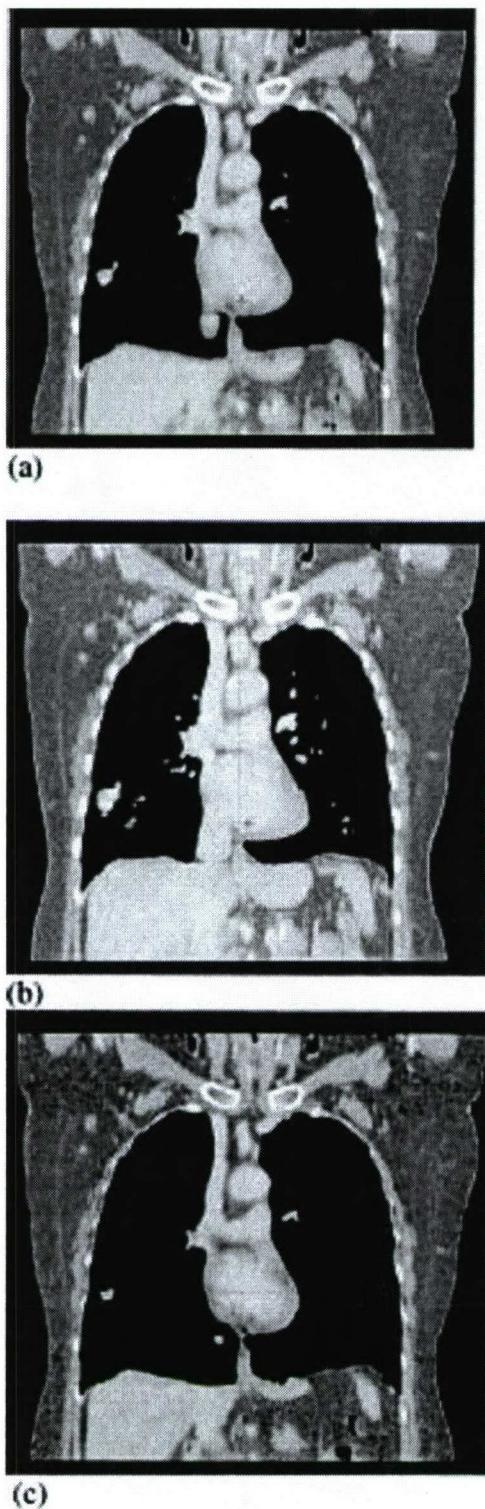


Fig. 6. Composite scans of a 4D CT lung patient. (a) Average pixel; (b) maximum-intensity pixel; (c) minimum-intensity pixel. The maximum-intensity pixel composite reveals the motion extent of hyperdense tissue (e.g., lung tumor), while the minimum-intensity pixel view provides the motion range of hypodense regions (e.g., lung air volume).

can proceed in a similar fashion as conventional 3D inverse planning to derive the optimal trajectories of the movements of the MLC leaves. An aperture-based optimization¹⁰⁷⁻¹⁰⁹ seems to be more adequate for dealing with the organ motion.⁹⁶

4D methods propose to track the tumor with the radiation beam as the tumor moves during the respiration cycle. These techniques require acquisition of some form of respiration signal (infrared reflective markers, spirometry, strain gauges, video tracking of chest outlines, and fluoroscopic tracking of implanted markers are some of the techniques employed to date), which is assumed to be correlated with internal anatomy motion. Fluoroscopy and the cine model electronic-portal-imaging device (EPID) have been proposed as a means for real-time guidance.^{110,113} While tumor tracking seems to be the ultimate goal of 4D radiation therapy, the real challenge is clarifying whether the 4D model is repeatable at the time of fractionated treatments, and determining how to correctly synchronize the MLC movements with the patient breathing. Real-time imaging and/or adaptive approaches will likely play a role in this aspect and the issue will surely need more research for many years to come.

INTER-FRACTION ORGAN MOVEMENT

Current techniques in dealing with inter-fraction organ movement

Uncertainty in patient setup has long been known as a limiting factor to conformal radiation therapy. Currently, the accuracy of patient setup is verified by megavoltage (MV) radiograph acquired with either radiographic film or EPID.^{114,115} The patient's bony landmarks are used to guide patient alignment. Poor soft tissue contrast and often unclear projection of the bony anatomy are major problems of the approach. To improve the situation, planar kV x-ray imaging has been implemented in a variety forms.¹¹⁰⁻¹¹³ While these systems show significantly increased contrast for bony structure differentiation, observing soft-tissue detail remains problematic and correction of daily organ motion is still challenging. Attempts have been made to use CT imaging to facilitate the patient setup process. Along this line, the offline adaptive-radiation-therapy (ART) strategy¹¹⁶ and in-room CT approach¹¹⁷ have been studied. The former method aims to partially compensate for organ motion by carrying out multiple CT scans in consecutive days in the first week of treatment. The image data are then employed to construct a patient specific PTV model from the composite of the CTVs with inclusion of statistical variations of the observed motions. While beneficial, the approach is hardly an ideal solution for dealing with the inter-fraction organ motion. It relies on establishing a statistical ensemble of all possible setup scenarios under a strong assumption that a limited number of off-line CT scans can ade-

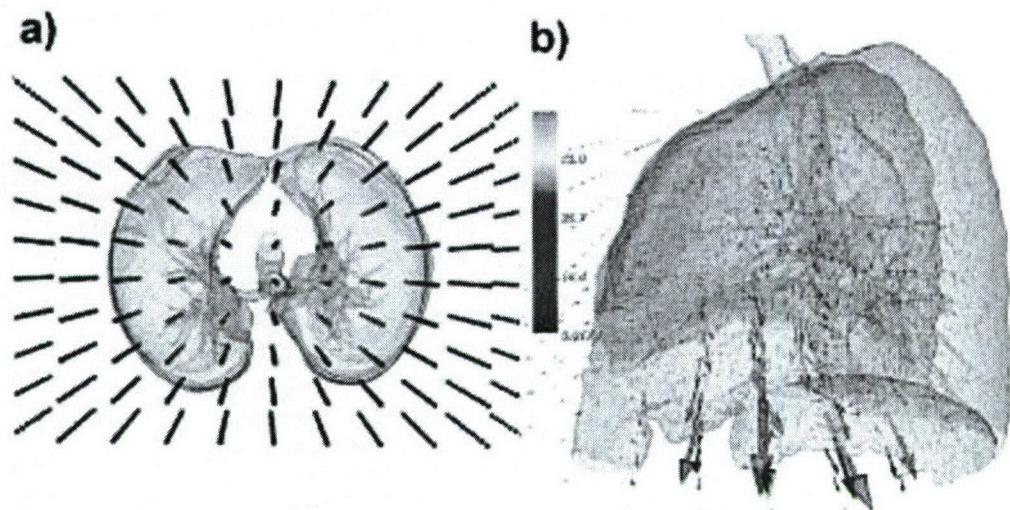


Fig. 7. (a) The BSpline grid superimposed on lung contours. (b) On each node, deformation is represented by arrows, where arrow length is proportional to the deformation.

quately describe the inherently complex, often unpredictable inter-fraction organ motion. Even when it is achievable, the ART margin is not optimized on a daily basis and there is still room for further improvement. An integrated CT/LINAC combination, in which the CT scanner is located inside the radiation therapy treatment room and the same patient couch is used for CT scanning and treatment (after a 180° couch rotation), should allow more accurate correction of interfractional setup errors. Some major radiotherapy vendors provide options to install a CT scanner in the treatment room. The overall precision of EXaCT Targeting™ from Varian has been evaluated by Court *et al.*¹¹⁷ However, the approach assumes a fixed relationship between the LINAC isocenter and the CT images and relies heavily on the mechanical integrity of the 2 otherwise independent systems. Increased capital cost and prolonged imaging and treatment are other concerns.

Other patient localization techniques available include ultrasound-based methods, video-based surface tracking, on-board cone-beam CT or kV x-ray imaging, CyberKnife and Tomotherapy, etc. For prostate radiation therapy, on-line ultrasound imaging has gained substantial interest^{118–120} but in practice has been found susceptible to subtle sources of error and inter-user variability. On-board CBCT holds promise to become a robust integrated on-line imaging technology that can yield unambiguous soft-tissue detail at the time of treatment. Furthermore, CT numbers correlate directly with electron density, thereby providing the potential for reconstruction of the actual dose delivered on a daily basis, in addition to simple anatomic structure alignment. The details of emerging CBCT will be presented in the next section. The robotic CyberKnife™ from Accuray Inc. (Sunnyvale, CA) represents another promising tech-

nology. The system has a feedback mechanism in which motion of the CTV, determined through the Accutrac infrared-x-ray-correlated imaging system, can be fed back to the robot to track the CTV.¹²¹ However, while this improves the duty cycle, there is a finite time between measuring tumor position and arranging the compensation for motion. Helical tomotherapy is an alternative means of delivering IMRT using a device that combines features of a linear accelerator and a helical CT scanner.¹²² The commercial version, the HI-ART IITM, can generate CT images using the same MV radiation beam that is used for treatment. Because the unit uses the actual treatment beam as the x-ray source for image acquisition, no surrogate telemetry systems are required to register image space to treatment space. Objective measures of noise, uniformity, contrast, linearity, and spatial resolution, and comparison with that of a commonly utilized CT simulator, have recently performed by Meeks *et al.*¹²³

CBCT for patient localization

CBCT based upon flat-panel technology integrated with a medical linear accelerator has recently become available from linac vendors for therapy guidance. The volumetric images may be used to verify and correct the planning patient setup in the linac coordinates by comparing with the patient position defined in treatment plan. Both kV and MV beams have been utilized for the application. The former typically consists of a kV-source and flat-panel combination mounted on the drum of a medical accelerator,¹²⁴ with the kV imaging axis orthogonal to that of MV therapy beam. The fan-beam and cone-beam MV CT in clinical applications have been reported by Meeks *et al.*¹²³ and Pouliot *et al.*¹²⁵, respectively. It appears that the MV images contain sufficient

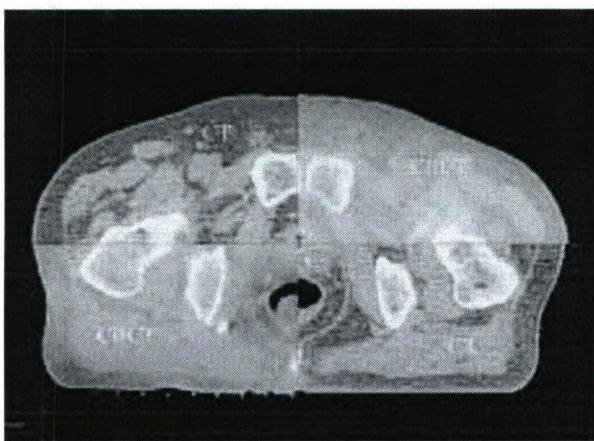


Fig. 8. The fusion of the 2 types of CT and CBCT images for a prostate case.

resolution of bone and air cavities to register them to structures in the planning CT with millimeter precision.^{124,125}

Currently, CBCT is primarily used for guiding the patient setup.^{126,127} The procedure is not much different from the current patient treatment, other than the fact that the AP/LAT portal images are replaced by volumetric data. In Fig. 8, we show 3D CBCT images of a prostate case in one of the fractional treatments fused with the patient's planning CT image. It is seen that soft-tissue structures and boundaries are visible to varying degrees in the CBCT images. The patient has implanted fiducials, which show up on both CBCT and planning CT. Our experience indicates that the cone beam data can clearly reveal setup error, as well as the anatomical deformations and other physiological changes. During the patient setup process, the 3D CBCT images are registered with the planning CT data through the use of either manual or automated 3D image registration software that calculates shifts in x-, y-, and z-directions (depending on the manufacturer, rotations can also be included). The movements determined during the registration represent the required setup corrections that should be applied to the patient. Both phantom and patient studies from our group have shown that the volumetric imaging is superior to the conventional MV or kV AP/LAT patient setup procedure. We note that, if only translational shifts are permissible, the level of improvement is generally within 2 mm as compared with kV AP/LAT setup procedure (2D/2D match). However, CBCT can readily detect rotational errors that may otherwise be missed by the 2D/2D match. In Fig. 9, we show the localization image for a head phantom with kV/kV 2D/2D match and 3D/3D match (CBCT/planning CT). The latter approach was found to be sensitive enough to identify a rotational error as small as 2°.

In practice, much effort is needed to improve the robustness and efficiency of the volumetric image regis-

tration process. Furthermore, to fully utilize the volumetric data, a new paradigm with seamlessly integrated simulation, planning, verification, and delivery procedure is urgently needed. Until this is realized clinically, the volumetric imaging is nothing but an expensive extension of the already functional planar verification approach. The capital cost and other related overheads do not seem to justify the marginal benefit if the volumetric data is simply used for determining the patient shift in the space. However, one should not underestimate the potential of the volumetric imaging for the future of radiation therapy, as it opens a new avenue (perhaps the only avenue), for us to realize the planned dose distribution with high confidence in clinical settings.

A few groups are working on deformable model based segmentation and patient setup procedures.^{93,126–128} When deformable registration is used, there are a few options to achieve the registration depending on whether the primary aim is to match soft-tissue, or to align 3D bony structures. In Fig. 10, we show a patient's CBCT and planning CT registration results using different registration schemes. The multiple choices result from the fact that the dimensionality of the patient data is much greater than that in the patient setup procedure, and suggest that deformable registration is not the ultimate solution to volumetric image-guided radiation therapy. Nonetheless, the technique improves the current method,¹²⁷ because it partially takes into account organ deformation by achieving the closest overlay match possible between the planning and CBCT data sets according to our clinical objective, and serves as a useful interim solution before a better integrated approach becomes available.

CBCT-based dose verification

Another important application of on-board volumetric imaging is verification of dose delivery. We have recently evaluated the accuracy of kV CBCT-based dose calculation and examined if current CBCT is suitable for the daily dose verification of patient treatment.^{129,130} A CT-calibration phantom was first used to calibrate both conventional CT and CBCT. CT and CBCT images of the calibration phantom, an anthropomorphic phantom and 2 patients (a lung and a prostate case) were then acquired for this study. Our results indicated that the imperfect quality of CBCT images has minimal impact (< 3%) on the dosimetric accuracy when the intra-fractional organ motion is small. When intra-fractional organ motion is large and motion artifacts are severe (e.g., in the case of lung cancer), the dosimetric discrepancy due to the poor image quality of current CBCT was found to be clinically significant. Furthermore, in the latter case, we found that it is possible to use a deformable registration algorithm to map the corresponding electron density information from planning CT to CBCT and then to proceed with conventional dose calculation.

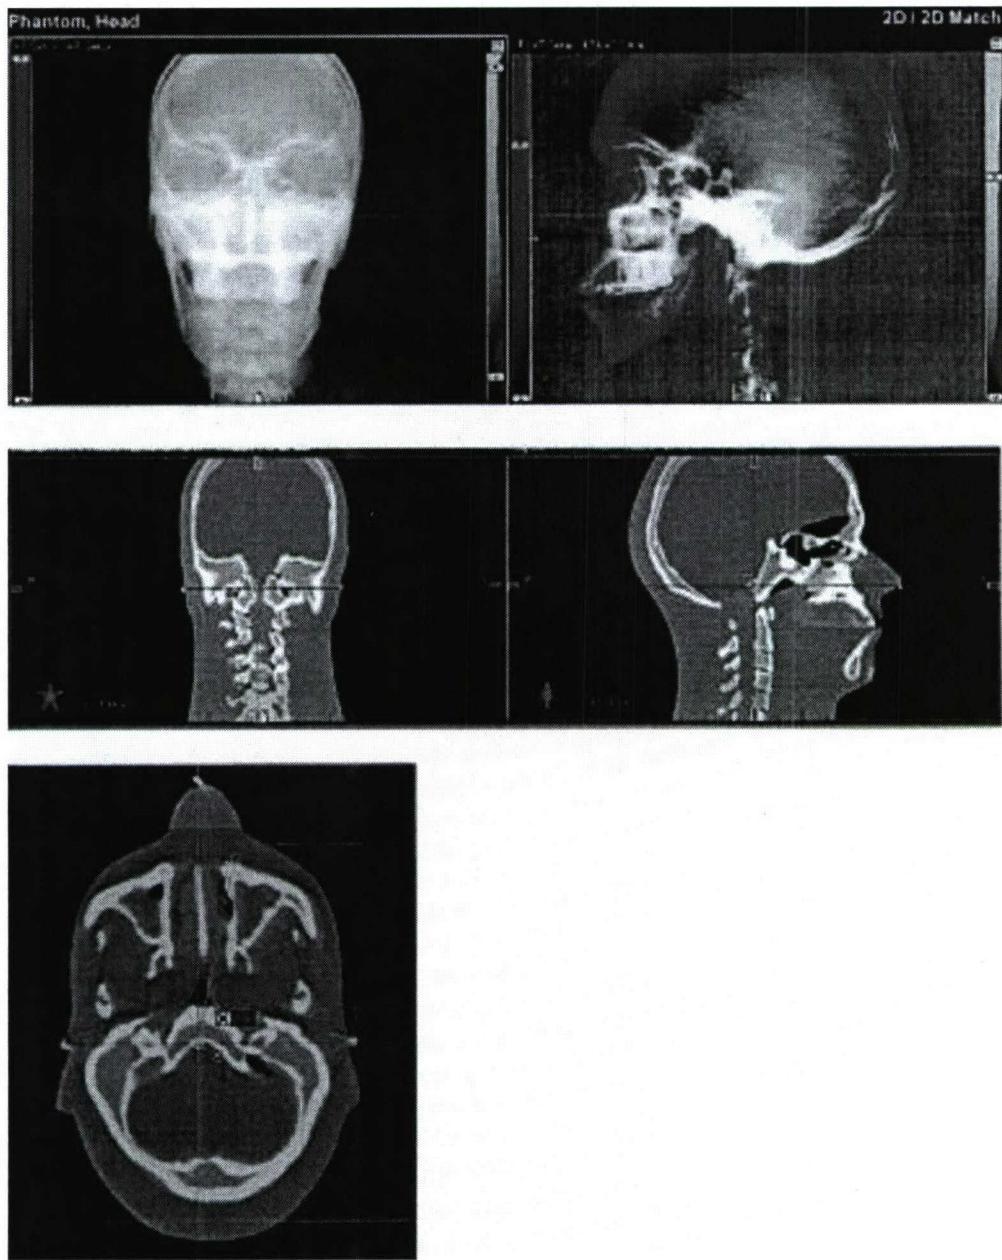


Fig. 9. Setup localization image for head phantom with kV/kV 2D/2D match (top) and 3D/3D CBCT match (middle). The image shown in the bottom panel illustrates that the CBCT is a sensitive technique capable of picking up a 2° rotational miss-match between the planning CT and CBCT.

Respiratory motion artifacts in CBCT

Superior to the common approaches based on 2 orthogonal images, CBCT can provide high-resolution 3D information of the patient in the treatment position, and thus has great potential for improved target localization and irradiation dose verification. In reality, however, scatter and organ motion are 2 major factors limiting the quality of current CBCT. When CBCT is used in imaging thorax or upper abdomen of a patient, respiration induced artifacts such as blurring, doubling, streaking, and distortion are observed, which heavily degrade the

image quality, and affect the target localization ability, as well as the accuracy of dose verification. These artifacts are much more severe than those found in conventional CT exams, in which each rotation of the scan can be completed within one second. On the contrary, in CBCT scan, the gantry rotation speed is much slower, typically 40 seconds to 1 minute for a full 360° scan in acquiring the projection data, which is more than 10 breathing cycles for most patients. In Fig. 11, we show the influence of the same motion on a regular “fast” CT scanner and CBCT for a motion phantom, where it is clearly seen

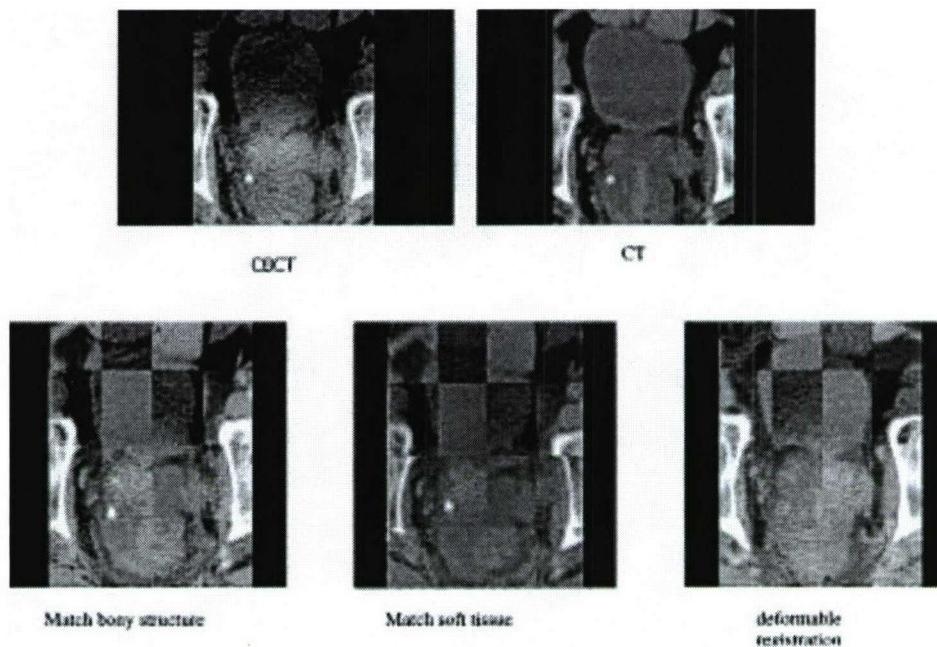


Fig. 10. Image registration of CBCT and planning CT based on bony structure matching, soft tissue matching, and deformable registration. Different matching techniques emphasize on different aspect of the multidimensional problem.

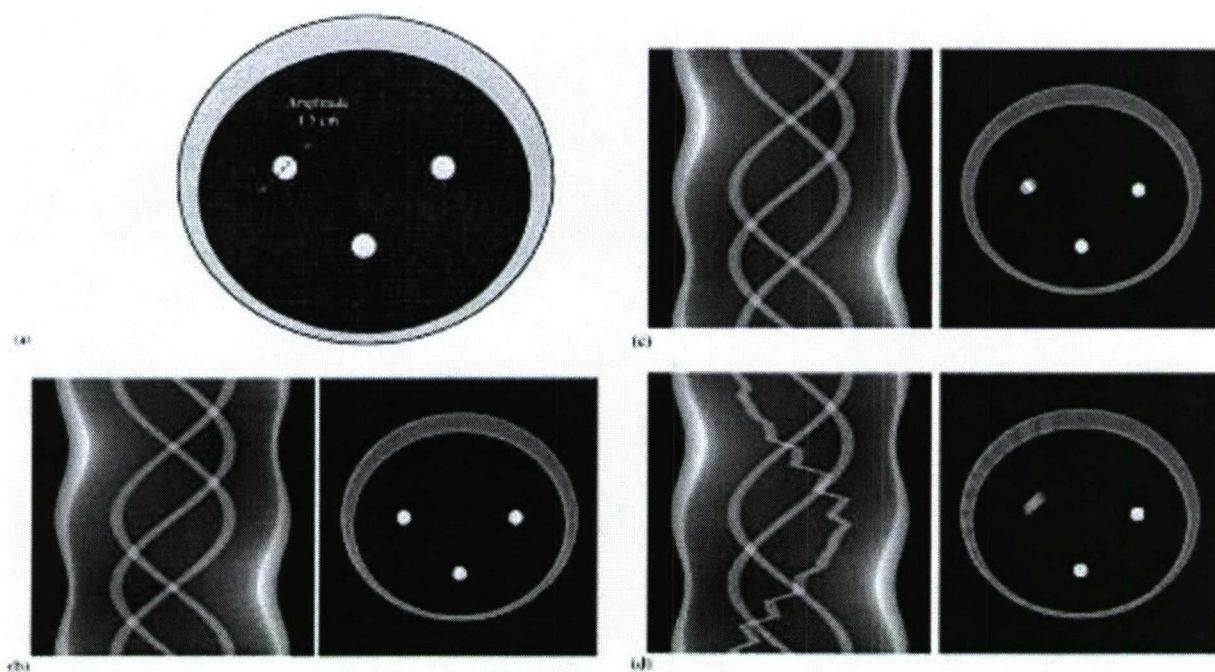


Fig. 11. (a) Motion phantom for CT and CBCT simulation study. The left circle moves diagonally with an amplitude of 1.5 cm and a period of 3.52 seconds. (b) Simulated sinograms and their corresponding reconstructed images with standard FBP algorithm when the circles are stationary. (c) and (d) show the sinograms and their corresponding reconstructed images for 1 s/rotation acquisition (conventional CT scan speed) and 40 s/rotation acquisition (on-board CBCT scan speed), respectively.

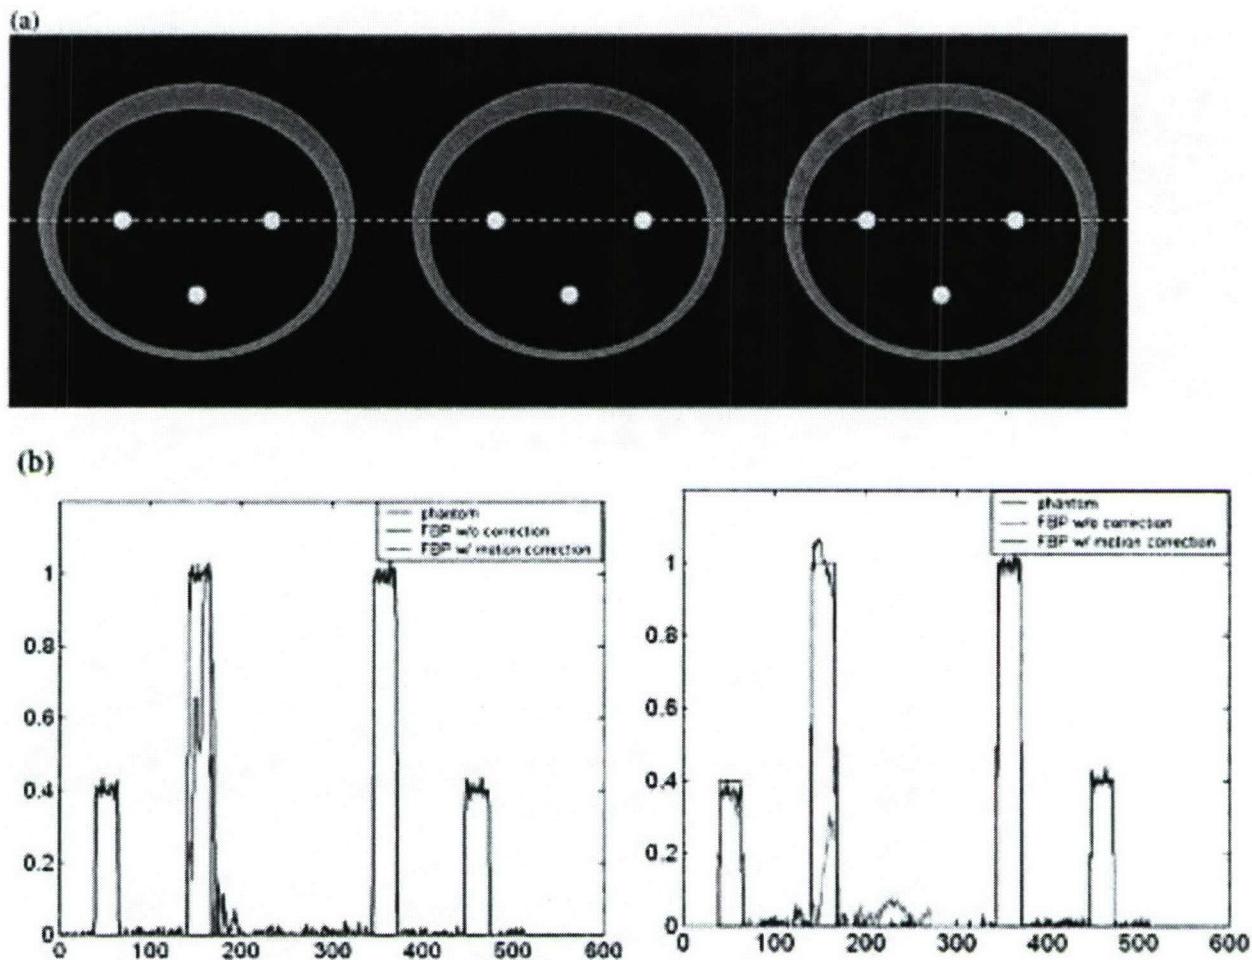


Fig. 12. (a) Phantom and images reconstructed with motion correction for CT and CBCT settings. The 3 images represent the reconstructed image of stationary phantom (left), the conventional “fast” CT (middle), and the CBCT (right). (b) Horizontal profiles through the moving circle for the reconstructed CT (left panel) and CBCT (right panel) images shown in the middle and right of (a). The profiles are in blue. For comparison, the profiles for the stationary phantom (left panel) and images reconstructed without motion artifacts removal mechanism are also plotted in each case (black and red curves, respectively).

that the motion artifacts are much greater than that in a fast scanner.

In the last decade, considerable effort has been devoted to finding solutions to remove motion artifacts and to obtain time-resolved medical images. Wang and Vannier¹³¹ presented a patient-motion estimation and compensation technique for helical CT systems. Willis and Bresler¹³² cast the motion artifact problem as a time-varying tomography problem and required special-purpose hardware to optimally sample the spatially and temporally band-limited CT signal space. A parametric model for the respiratory motion was used in MRI, and the motion artifacts were successfully reduced by modifying the reconstruction algorithm.¹³³ Crawford *et al.*¹³⁴ brought the concept into CT imaging, and derived an exact reconstruction formula for motion compensation for CT scans. Generally, motion correction algorithms that assume a motion model work well when the motion

conforms to the model, but have limited success when it does not. As described above, 4D CT has been developed in radiation oncology application to explicitly account for the respiratory motion. The 4D CT can be used to derive a patient-specific deformation field and then incorporated into the CBCT filtered-backprojection (FBP) image reconstruction process.¹³⁵ The algorithm was tested with simulations at different settings corresponding to conventional CT and CBCT scan protocols, with translational motion and more complex motion, and with and without Gaussian noise. In Fig. 12, we show the result for the motion phantom depicted in Fig. 11.¹³⁵ Because the motion model is directly derived from the patient images, it should be more accurate than other artificial modeling, and therefore more efficient motion correction is expected. In addition to this approach, Sonke *et al.*¹³⁶ developed a CBCT procedure consisting of retrospective sorting in projection space, similar to

that used in 4D CT. The subsets of projection data are then reconstructed into 4D CBCT dataset. To achieve a sufficient temporal resolution, however, this will require slowing down the gantry rotation. The assumption of periodicity of the respiratory motion is also necessary. Li *et al.*¹³⁷ have recently established a novel 4D CBCT reconstruction formalism, in which the reconstruction of a phase will consider not only the projections corresponding to that phase but also those of other phases. By incorporating information from other phases, the efficacy and quality of 4D CBCT images are substantially improved. Zeng *et al.*¹³⁸ proposed a method to estimate the parameters of a non-rigid, free-breathing motion model from a set of projections of thorax that are acquired using a slow rotating CBCT scanner.

RIGID AND DEFORMABLE IMAGE REGISTRATION

Development of an effective image registration technique has been one of the most important research areas. Depending on the mathematical nature of the transformation, image registration is divided into rigid and deformable registrations. In rigid transformations, it is assumed that the geometry of the object is identical in the 2 input images and no distortion occurs in the image acquisition process. A rigid transformation consists of 6 degrees of freedom: 3 displacement parameters and 3 rotational parameters. Deformable registration, on the other hand, is more complicated and entails the modeling of voxel dependent distortion. Clinically, the need for a robust image registration algorithm to compare/fuse images representing the same structures imaged under different conditions or on different modalities is ever increasing because of the extensive use of multi-modality imaging and the emergence of new imaging techniques and methods.

Computer-based rigid image registration has gained widespread popularity in the last decade and is used in routine clinical practice. In this approach, the matching of the 2 input images is formulated into an optimization problem and the best registration of the 2 images is obtained by iteratively comparing various possible matches until no better registration can be found. The search for the optimal match of the 2 input images is usually gauged by a ranking function constructed based on some physical considerations. Depending on the nature of the input images, the formulation of the problem can be highly complicated. Court and Dong¹³⁹ used a rigid transformation for the correction of tissue displacement. A deformable procedure based on the finite element model (FEM), in which images are described as blocks of elastic materials on which forces apply, was proposed by Bharath *et al.*¹⁴⁰ and Brock *et al.*¹⁴¹ In this approach, the parameters that control the behavior of the elastic material and are responsible for the conversion of forces into local deformations of the elastic material are

Young's elastic modulus and Poisson's ratio. Although powerful, the model has the drawback that values of the elasticity and density constant for various tissues are not readily available and have to be found by a trial and error procedure. The method also relies on using complicated software to generate a FEM mesh and masks of the involved structures. Schreibmann and Xing have proposed a general narrow-band approach for deformable registration.⁹³ Depending on the problem, modeling of individual voxel movement can also be made using either B-splines,⁷⁹ thinplate splines^{142,143}, optical flow algorithms,¹⁴⁴ or fluid flow algorithms.¹⁴⁵ Spline interpolation is a relatively simple approach and our experience with the algorithm suggested that the free-form registration is stable and accurate for dealing with IGRT image registration problems.¹⁴⁶ An improvement to this method can be achieved by using a spline model with the smoothness of the deformation field assured by the interpolation between a grid of fixed control points. A simple method along this line is to deduce the spline coefficients from a set of user-defined control points, as was done by Fei *et al.*¹⁴⁷ and Lian *et al.*¹⁴³ in warping and registration of MR volumes. Coselman *et al.*¹⁴⁸ used a similar technique to study the accuracy of mutual-information-based CT registration of the lung at exhale and inhale respiratory states.

To facilitate the computer decision-making process, image pre-processing or user interaction may be required, especially when dealing with a deformable image registration. The use of homologous anatomic landmark pairs on the 2 input images or the control points is an example of this. In reality, the user must have a detailed understanding of the patient anatomy and the characteristics of the 2 modalities in order to accurately identify the control points on both images. The point pairs are usually obtained interactively with the user repetitively exploring the input image sets and each time trying to locate a point in both of them. Due to the 3D nature, the process is rather tedious and difficult to perform. Schreibmann and Xing¹⁴⁹ have developed a general method to facilitate the selection of control points for both rigid and deformable image registrations. Instead of relying on the interactive selection of homologous control point pairs on both model and reference images, in the proposed approach the user needs only to identify some small control volumes on the model image in a somewhat arbitrary fashion. This new way of image registration eliminates the need for the manual placement of the homologous control points and allows us to register the 2 images accurately and efficiently. The method was applied to both rigid and non-rigid image registration problems and our results indicated that the registration is reliable and provides a valuable tool for intra- or inter-modality image registration. In Fig. 13, we show the registration result of a rectal cancer patient who has undergone both CT and FLT-PET scans. The increased robustness and confidence in the registration and the

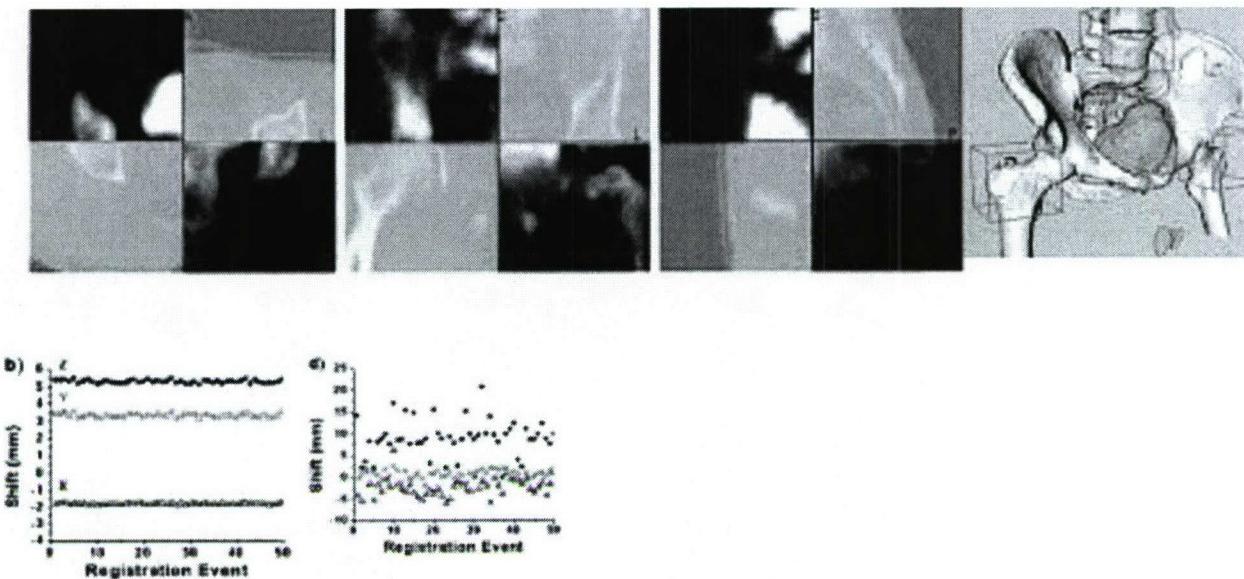


Fig. 13. Sagittal, coronal, and axial views of the CT and FLT-PET registration. In addition to the checkerboard display, a 3D view of the registration is also presented, where an excellent coincidence is observed between the bony structures revealed in CT (white) and PET images (orange). The right 2 panels of the 2nd row show the convergence behaviors of our method and the conventional method for 50 independent calculations. Our method leads to reproducible shifts in x-, y-, and z-directions, and the conventional approach based on the information contained in the whole image entity leads to large variations in the shifts.

increased speed of calculation, especially in the case of the deformable registration, are important features of the new technique. Compared to the manual rigid registration, this method eliminates the nuisance of the control point pair selection and removes a potential source of error in registration. Compared to the automated method, the technique is more intuitive and robust, especially in the presence of image artifacts.

CLINICAL EXPERIENCE WITH IGRT

Clinically implemented IGRT techniques at Stanford include 4D CT, 4D PET, Varian OBI (both planar and CBCT), gating, and Accuray CyberKnife. Several image-guided clinical protocols are under investigation. 4D CT/PET information are used in about 40% of the thorax and upper abdomen cases for patient specific tumor margin definition in 3.5D radiation therapy or for treatment planning of gated radiation therapy. CBCT is mainly applied for patient setup in the treatment of head-and-neck, and prostate and other pelvic diseases. For these sites, the CBCT image quality is reasonable to visualize soft tissues, but the quality is generally notably inferior to that of the state-of-the-art multi-slice fan beam CT scanner. Scan truncation artifacts because the patient shadow does not fit on the detector and/or organ motion often cause Hounsfield unit calibration problems. While this does not seem to influence the image registration, the use of CBCT for dose calculation should proceed with caution. Our initial experience indicates that, when com-

pared with traditional CT-based calculation, the dosimetric error is typically less than 3% for prostate or head-and-neck cases but could be significantly greater in the thoracic region. Comparison between cone beam data and portal image derived setup errors show only slight differences (< 2 mm). However, we should note that the differences are derived purely based on the use of manufacturer-provided image-fusion software, which often emphasizes the high-intensity voxels in bony structures. The next step is to implement soft-tissue based setup corrections clinically. In reality, volumetric data contain much more information compared to planar images, and CBCT promises to be more useful in the future when it is better integrated with treatment planning and delivery systems. An ideal integration would be to use volumetric image-derived information to “tweak” or re-optimize the treatment plan. This work is still in progress at Stanford.

As another example of IGRT treatment, we describe our phase I and II pancreatic tumor dose escalation protocol. The aim is to use CyberKnife to target pancreatic tumors more precisely and to limit the toxicities associated with treatment. In a phase I study, we treated patients with a single fraction of 15, 20, and 25 Gy to unresectable pancreatic tumors using the Cyberknife stereotactic radiotherapy (SRT) system (Accuray).¹⁵⁰ To track tumor movement, we implant fiducial seeds percutaneously into the pancreatic tumor. Using the Accuray Synchrony platform, a model in which the position of the internal fiducials is correlated with the patient’s respi-

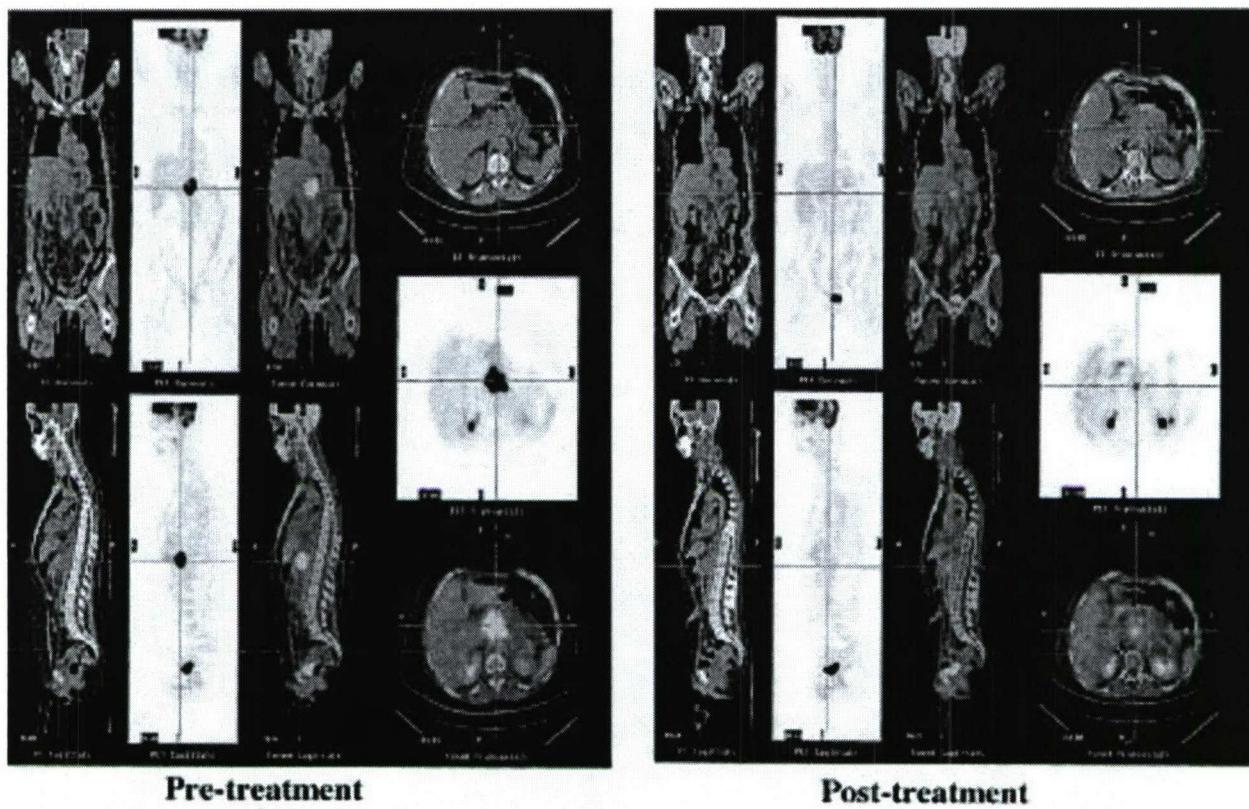


Fig. 14. FDG-PET images of a pancreatic patient before and after radiation therapy.

tory motion is developed. The Cyberknife is able to make real-time corrections to compensate for tumor movement during respiration. Prior to treatment, patients underwent 4D planning CT scans. Using this data set, we are able to visualize how the pancreatic tumor moves/deforms through respiration and compensate for these dynamic changes.¹⁵¹ Minimal acute gastrointestinal toxicity was observed even at the highest dose. All patients who received 25 Gy had no further local progression of their tumor until death. In a follow up phase II study, a cohort of 19 patients were treated with 45-Gy conventionally fractionated radiation therapy using IMRT to the pancreas and regional lymph nodes followed by a 25-Gy Cyberknife stereotactic radiotherapy boost to the primary tumor.¹⁵² An excellent rate of local control with this therapy was confirmed. Because of the rapid progression of systemic disease, we did not observe a significant improvement in overall survival as compared to historic controls. However, most patients had a clinical benefit (decreased pain, increased activity) and decreased serum tumor marker for pancreatic cancer (CA-19-9) following therapy. To document that SRS truly resulted in an anti-tumor effect, we routinely obtain FDG-PET/CT scans before and after treatment. Figure 14 is an example of one such study. There was intense metabolic activity of the pancreatic tumor prior to therapy with a near complete resolution of FDG uptake in this patient 4

weeks following therapy. The technological challenge for IGRT to minimize toxicity in this clinical scenario is the precision delivery of high-dose radiotherapy. This cannot be accomplished without taking into account the respiratory associated motion of pancreatic tumors. This movement takes place in multiple planes and can result in tumor displacement of up to 1–2 cm. Furthermore, tumor and organ deformation during respiration must also be compensated for during therapy.

SUMMARY

With the development of IMRT during the 1990s, radiation therapy entered a new era. This new process of treatment planning and delivery shows significant potential for improving the therapeutic ratio and offers a valuable tool for dose escalation and/or radiation toxicity reduction. The improved dose conformity and steep dose gradients necessitate enhanced precision and accuracy in patient localization and spawn the development of IGRT, in which various metabolic and anatomical imaging techniques are integrated into the radiation therapy process. The overall goal of IGRT is to target tumors more accurately while better sparing the normal tissues. Much recent effort is focused on removing the uncertainty in the definition of the target volume and in the determination of the position of mobile and often deformable

organs. Biological imaging described in this article will allow us not only to delineate the boundary of the tumor volume based on the tumors' biological characteristics but also to map out the biology distribution of the cancer cells, affording a significant opportunity for BCRT treatment in the future. Developments of effective 4D CT/PET techniques will provide effective means for us to understand the temporal dependence of the involved structures and design the best possible strategy for targeting the moving tumor. Integration of various imaging tools for off-line and on-line application is also of paramount importance, enabling us to ensure the planned dose distributions can be realized in the clinical setting. With the newly available IGRT tools, physicians will be able to optimize radiotherapy accuracy and precision by adjusting the radiation beam based on the actual positions of the target tumor and critical organs during radiation therapy planning and treatment. We should mentioned that IGRT is still in its infancy and many technical issues remain to be resolved, such as the establishment of a robust deformable registration method, auto-mapping of the contours outlined on the planning CT to CBCT or to different phases of 4D CT, and management of the sheer volume of acquired image sets (both 4D CT/PET and CBCT). However, it is believed that much of these technical hurdles will be resolved with time, and that IGRT will become the standard of practice in the future through the effort of researchers around the world.

Acknowledgments—The authors thank the following individuals for their input and help: Drs. S. Hancock, C. King, K. Goodman, Q. Le, B. Loo, K. Horse, A. Boyer, D., P. Maxim, A. Hsu, T. Pawlicki, B. Ambroster, Spielman, S. Gambhir, A. Quon, D. Levy from Stanford University; S. Johnson, R. Morse, P. Munro, C. Zankowski, R. Wicha, M. Svatos, K. Kennedy, and R. Stark from Varian Medical Systems; G. Chen from Massachusetts General Hospital; and G. Mageras from Memorial Sloan Kettering Cancer Center. This work is supported in part by a research grants from the National Cancer Institute (5R01 CA98523-01), Department of Defense (PC040282), the American Cancer Society (RSG-01-022-01-CCE), and Varian Medical Systems. Permission for using copyrighted figures from *Medical Physics* and *International Journal of Radiation Oncology, Biology, Physics* is gratefully acknowledged.

REFERENCES

- IMRT Collaborative Working Group: Intensity-modulated radiotherapy: Current status and issues of interest. *Int. J. Radiat. Oncol. Biol. Phys.* **51**:880–914; 2001.
- AAPM IMRT Sub-Committee: Guidance document on delivery, treatment planning, and clinical implementation of IMRT: Report of the IMRT Subcommittee of the AAPM Radiation Therapy Committee. *Med. Phys.* **30**:2089–115; 2003.
- Galvin, J.M.; Ezzell, G.; Eisbrauch, A.; *et al.* Implementing IMRT in clinical practice: A joint document of the American Society for Therapeutic Radiology and Oncology and the American Association of Physicists in Medicine. *Int. J. Radiat. Oncol. Biol. Phys.* **58**:1616–34; 2004.
- Hamilton, R.J.; Sweeney, P.J.; Pelizzari, C.A.; *et al.* Functional imaging in treatment planning of brain lesions. *Int. J. Radiat. Oncol. Biol. Phys.* **37**:181–8; 1997.
- Sawada, A.; Yoda, K.; Kokubo, M.; *et al.* A technique for noninvasive respiratory gated radiation treatment system based on a real time 3D ultrasound image correlation: A phantom study. *Med. Phys.* **31**:245–50; 2004.
- Yu, Y.; Anderson, L.L.; Li, Z.; *et al.* Permanent prostate seed implant brachytherapy: Report of the American Association of Physicists in Medicine Task Group No. 64. *Med. Phys.* **26**:2054–76; 1999.
- Pirzkall, A.; McKnight, T.R.; Graves, E.E.; *et al.* MR-spectroscopy guided target delineation for high-grade gliomas. *Int. J. Radiat. Oncol. Biol. Phys.* **50**:915–28; 2001.
- Kurhanewicz, J.; Hricak, H.; Sokolov, D.L.; *et al.* Prostate cancer: Localization with three-dimensional proton MR spectroscopic imaging—clinicopathologic study. *Radiology* **206**:785–90; 1998.
- DiBiase, S.J.; Hosseinzadeh, K.; Gullapalli, R.P.; *et al.* Magnetic resonance spectroscopic imaging-guided brachytherapy for localized prostate cancer. *Int. J. Radiat. Oncol. Biol. Phys.* **52**:429–38; 2002.
- Kim, D.H.; Margolis, D.; Xing, L.; *et al.* In vivo prostate magnetic resonance spectroscopic imaging using two-dimensional J-resolved PRESS at 3 T. *Magn. Reson. Med.* **53**:1177–82; 2005.
- Zaider, M.; Zelefsky, M.J.; Lee, E.K.; *et al.* Treatment planning for prostate implants using magnetic-resonance spectroscopy imaging. *Int. J. Radiat. Oncol. Biol. Phys.* **47**:1085–96; 2000.
- Hunjan, S.; Gibbs, I.; Spielman, D.M.; *et al.* Validating magnetic resonance spectroscopic imaging for radiation therapy guidance. *Int. J. Radiat. Oncol. Biol. Phys.* **57**:1159–73; 2003.
- Townsend, D.W.; Carney, J.P.; Yap, J.T.; *et al.* PET/CT today and tomorrow. *J. Nucl. Med.* **45**:4S–14S; 2004.
- Czernin, J.; Schelbert, H. PET/CT imaging, facts, options, hopes, and questions. *J. Nucl. Med.* **45**:1S–3S; 2004.
- Hicks, R.J.; Kalff, V.; MacManus, M.P.; *et al.* (18)F-FDG PET provides high-impact and powerful prognostic stratification in staging newly diagnosed non-small cell lung cancer. *J. Nucl. Med.* **42**:1596–604; 2001.
- Mac Manus, M.P.; Hicks, R.J.; Ball, D.L.; *et al.* F-18 fluorodeoxyglucose positron emission tomography staging in radical radiotherapy candidates with nonsmall cell lung carcinoma: Powerful correlation with survival and high impact on treatment. *Cancer* **92**:886–95; 2001.
- Bradley, J.D.; Dehdashhti, F.; Mintun, M.A.; *et al.* Positron emission tomography in limited-stage small-cell lung cancer: A prospective study. *J. Clin. Oncol.* **22**:3248–54; 2004.
- Howard, A.; Mehta, M.P.; Ritter, M.A.; *et al.* The value of PET/CT in gross tumor volume delineation in lung and esophagus cancer. *Int. J. Radiat. Oncol. Biol. Phys.* **60**:S536–537; 2004.
- Hofer, C.; Laubenbacher, C.; Block, T.; *et al.* Fluorine-18-fluorodeoxyglucose PET is useless for the detection of local recurrence after radical prostatectomy. *Eur. Urol.* **36**:31–5; 1999.
- Liu, I.J.; Zafar, M.B.; Lai, Y.H.; *et al.* Fluorodeoxyglucose PET studies in diagnosis and staging of clinically organ-confined prostate cancer. *Urology* **57**:108–11; 2001.
- Gambhir, S.S.; Srinivasan, A.; Banerjee, P.K.; *et al.* PET in oncology: Will it replace the other modalities? *J. Nucl. Med.* **39**:729–34; 1998.
- Gambhir, S.S. Molecular imaging of cancer with PET. *Nat. Rev. Cancer* **2**:683–93; 2002.
- Schöder, H.; Larson, S.M. Positron emission tomography for prostate, bladder, and renal cancer. *Semin. Nucl. Med.* **34**:274–92; 2004.
- Buck, A.K.; Halter, G.; Schirrmeyer, H.; *et al.* Imaging proliferation in lung tumors with PET: 18F-FLT versus 18F-FDG. *J. Nucl. Med.* **44**:1426–31; 2003.
- Shields, A.F.; Grierson, J.R.; Dohmen, B.M.; *et al.* Imaging proliferation in vivo with [F-18]FLT and positron emission tomography. *Nat. Med.* **4**:1334–6; 1998.
- Buck, A.K.; Schirrmeyer, H.; Hetzel, M.; *et al.* 3-deoxy-3-[(18)F]fluorothymidine-PET for noninvasive assessment of proliferation in pulmonary nodules. *Cancer Res.* **62**:3331–4; 2002.
- Smyczek-Gargya, B.; Fersis, N.; Dittmann, H.; *et al.* PET with [18F]fluorothymidine for imaging of primary breast cancer: A pilot study. *Eur. J. Nucl. Med. Mol. Imag.* **31**:720–724.
- Shreve, P.D.; Gross, M.D. Imaging of the pancreas and related diseases with PET carbon-11-acetate. *J. Nucl. Med.* **38**:1305–10; 1997.

29. Seltzer, M.A.; Jahan, S.A.; Sparks, R.; *et al.* Radiation dose estimates in humans for (11)C-acetate whole-body PET. *J. Nucl. Med.* **45**:1233–6; 2004.
30. Jager, P.L.; Que, T.H.; Vaalburg, W.; *et al.* Carbon-11 choline or FDG-PET for staging of oesophageal cancer? *Eur. J. Nucl. Med.* **28**:1845–9; 2001.
31. Ishiwata, K.; Kasahara, C.; Hatano, K.; *et al.* Carbon-11 labeled ethionine and propionine as tumor detecting agents. *Ann. Nucl. Med.* **1**:115–22; 1997.
32. Iozzo, P.; Osman, S.; Glaser, M.; *et al.* In vivo imaging of insulin receptors by PET: Preclinical evaluation of iodine-125 and iodine-124 labelled human insulin. *Nucl. Med. Biol.* **29**:73–82; 2002.
33. Herlin, G.; Persson, B.; Bergström, M.; *et al.* 11C-harmine as a potential PET tracer for ductal pancreas cancer: In vitro studies. *Eur. Radiol.* **13**:729–3; 2003.
34. Brown, W.D.; Oakes, T.R.; DeJesus, O.T.; *et al.* Fluorine-18-fluoro-L-DOPA dosimetry with carbidopa pretreatment. *J. Nucl. Med.* **39**:1884–91; 1998.
35. Blankenberg, F.G.; Strauss, H.W. Nuclear medicine applications in molecular imaging. *J. Magn. Res. Imaging* **16**:352–61; 2002.
36. Oyama, N.; Akino, H.; Kanamaru, H.; *et al.* 11C-acetate PET imaging of prostate cancer. *J. Nucl. Med.* **43**:181–6; 2002.
37. Kotzerke, J.; Volkmer, B.G.; Glatting, G.; *et al.* Intraindividual comparison of [11]C-acetate and [11]C-choline PET for detection of metastases of prostate cancer. *Nuklearmedizin* **42**:25–30; 2003.
38. Mathews, D.; Oz, O.K. PET in prostate and renal cell carcinoma. *Curr. Opin. Urol.* **12**:381–5; 2002.
39. DeGrado, T.R.; Baldwin, S.W.; Wang, S.; *et al.* Synthesis and evaluation of (18)F-labeled choline analogs as oncologic PET tracers. *J. Nucl. Med.* **42**:1805–14; 2001.
40. Hara, T.; Kosaka, N.; Kishi, H. Development of (18)F-fluoroethylcholine for cancer imaging with PET: Synthesis, biochemistry, and prostate cancer imaging. *J. Nucl. Med.* **43**:187–99; 2002.
41. Hara, T.; Kosaka, N.; Kishi, H. PET imaging of prostate cancer using carbon-11-choline. *J. Nucl. Med.* **39**:990–5; 1998.
42. Sutinen, E.; Nurmi M.; Roivainen A.; *et al.* Kinetics of [(11)C]choline uptake in prostate cancer: A PET study. *Eur. J. Nucl. Med. Mol. Imaging* **31**:317–24; 2004.
43. Chen, X.; Park, R.; Hou, Y.; *et al.* MicroPET and autoradiographic imaging of GRP receptor expression with 64Cu-DOTA-[Lys3]bombesin in human prostate adenocarcinoma xenografts. *J. Nucl. Med.* **45**:1390–7; 2004.
44. Rasey, J.S.; Koh, W.J.; Evans, M.L.; *et al.* Quantifying regional hypoxia in human tumors with positron emission tomography of [18F]fluoromisonidazole: A pretherapy study of 37 patients. *Int. J. Radiat. Oncol. Biol. Phys.* **36**:417–28; 1996.
45. O'Donoghue, J.A.; Zanzonico, P.; Pugachev, A.; *et al.* Assessment of regional tumor hypoxia using 18F-fluoromisonidazole and 64Cu(II)-diacetyl-bis(N4-methylthiosemicarbazone) positron emission tomography: Comparative study featuring microPET imaging, Po2 probe measurement, autoradiography, and fluorescent microscopy in the R3327-AT and FaDu rat tumor models. *Int. J. Radiat. Oncol. Biol. Phys.* **61**:1493–502; 2005.
46. Piert, M.; Machulla, H.J.; Picchio, M.; *et al.* Hypoxia-specific tumor imaging with 18F-fluoroazomycin arabinoside. *J. Nucl. Med.* **46**:106–13; 2005.
47. Chao, K.S.; Mutic, S.; Gerber, R.L.; *et al.* A novel approach to overcome hypoxic tumor resistance: Cu-ATSM-guided intensity-modulated radiation therapy. *Int. J. Radiat. Oncol. Biol. Phys.* **49**:1171–82; 2001.
48. de Jong I.J.; Pruijm, J.; Elsinga, P.H.; *et al.* 11C-choline positron emission tomography for the evaluation after treatment of localized prostate cancer. *Eur. Urol.* **44**:327–8; discussion 38–9; 2003.
49. Ling, C.C.; Humm, J.; Larson, S.; *et al.* Towards multidimensional radiotherapy (MD-CRT): Biological imaging and biological conformality. *Int. J. Radiat. Oncol. Biol. Phys.* **47**:551–60; 2000.
50. Brahme, A. Individualizing cancer treatment: Biological optimization models in treatment planning and delivery. *Int. J. Radiat. Oncol. Biol. Phys.* **49**:327–37; 2001.
51. Xing, L.; Cotrutz, C.; Hunjan, S.; *et al.* Inverse planning for functional image-guided IMRT. *Phys. Med. Biol.* **47**:3567–78; 2002.
52. Alber, M.; Paulsen, F.; Eschman, S.M.; *et al.* On biologically conformal boost dose optimization. *Phys. Med. Biol.* **48**:N31–5; 2003.
53. Yang, Y.; Xing, L. Towards biologically conformal radiation therapy (BCRT): Selective IMRT dose escalation under the guidance of spatial biology distribution. *Med. Phys.* **32**:1473–84; 2005.
54. Seppenwoeld, Y.; Shirato, H.; Kitamura, K.; *et al.* Precise and real-time measurement of 3D tumor motion in lung due to breathing and heartbeat, measured during radiotherapy. *Int. J. Radiat. Oncol. Biol. Phys.* **53**:822; 2002.
55. Shirato, H.; Seppenwoeld, Y.; Kitamura, K.; *et al.* Intrafractional tumor motion: Lung and liver. *Semin. Radiat. Oncol.* **14**:10–8; 2004.
56. Shimizu, S.; Shirato, H.; Ogura, S.; *et al.* Detection of lung tumor movement in real-time tumor-tracking radiotherapy. *Int. J. Radiat. Oncol. Biol. Phys.* **51**:304–10; 2001.
57. Xu, Q.; Hamilton, R. Novel respiratory gating method based on automated analysis of ultrasonic diaphragm motion. *Med. Phys.* **32**:2124; 2005.
58. Huang, M.H.; Lin, Y.S.; Lee, C.L.; *et al.* Use of ultrasound to increase effectiveness of isokinetic exercise for knee osteoarthritis. *Arch. Phys. Med. Rehabil.* **86**:1545–51; 2005.
59. Seiler, P.G.; Blattmann, H.; Kirsch, S.; *et al.* A novel tracking technique for the continuous precise measurement of tumour positions in conformal radiotherapy. *Phys. Med. Biol.* **45**:N103–10; 2000.
60. Balter, J.M.; Wright, J.N.; Newell, L.J.; *et al.* Accuracy of a wireless localization system for radiotherapy. *Int. J. Radiat. Oncol. Biol. Phys.* **61**:933–7; 2005.
61. Nehmeh, S.A.; Erdi, Y.E.; Pan, T.; *et al.* Four-dimensional (4D) PET/CT imaging of the thorax. *Med. Phys.* **31**:3179; 2004.
62. Thorndyke, B.; Schreibmann, E.; Maxim, P.; *et al.* Enhancing 4D PET through retrospective stacking. Presented at 2005 Annual Meeting of American Association of Physicists in Medicine, Seattle, WA, 2005.
63. Thorndyke, B.; Schreibmann, E.; Xing, L. Reducing motion artifacts in PET through retrospective stacking. *Med. Phys.* **33**:2066; In press.
64. Vedam, S.S.; Keall, P.J.; Kini, V.R.; *et al.* Acquiring a four-dimensional computed tomography dataset using an external respiratory signal. *Phys. Med. Biol.* **48**:45; 2003.
65. Keall, P.J.; Stark, G.; Shukla, H.; *et al.* Acquiring 4D thoracic CT scans using a multislice helical method. *Phys. Med. Biol.* **49**:2053–67; 2004.
66. Low, D.A.; Nystrom, M.; Kalinin, E.; *et al.* A method for the reconstruction of four-dimensional synchronized CT scans acquired during free breathing. *Med. Phys.*
67. Rietzel, E.; Pan, T.; Chen, G.T.Y. Four-dimensional computed tomography: Image formation and clinical protocol. *Med. Phys.* **32**:874; 2005.
68. Li, T.; Schreibmann, E.; Thorndyke, B.; *et al.* Radiation dose reduction in 4D CT. *Med. Phys.* In press.
69. Kubo, H.D.; Len, P.M.; Minohara, S.; *et al.* Breathing-synchronized radiotherapy program at the University of California Davis Cancer Center. *Med. Phys.* **27**:346; 2000.
70. Rueckert, D.; Sonoda, L.I.; Hayes, C.; *et al.* Nonrigid registration using free-form deformations: Application to breast MR images. *IEEE Trans. Med. Imaging* **18**:712–21; 1999.
71. Pan, T.; Lee, T.-Y.; Rietzel, E.; *et al.* 4D-CT imaging of a volume influenced by respiratory motion on multi-slice CT. *Med. Phys.* **31**:333; 2004.
72. Kini, V.R.; Vedam, S.S.; Keall, P.J.; *et al.* Patient training in respiratory-gated radiotherapy. *Med. Dosim.* **28**:7; 2003.
73. Wong, J.W.; Sharpe, M.B.; Jaffray, D.A.; *et al.* The use of active breathing control (ABC) to reduce margin for breathing motion. *Int. J. Radiat. Oncol. Biol. Phys.* **44**:911; 1999.
74. Frazier, R.C.; Vicini, F.A.; Sharpe, M.B.; *et al.* Impact of breathing motion on whole breast radiotherapy: A dosimetric analysis using active breathing control. *Int. J. Radiat. Oncol. Biol. Phys.* **58**:1041–7; 2004.
75. Dawson, L.A.; Brock, K.K.; Kazanjian, S.; *et al.* The reproducibility of organ position using active breathing control (ABC) during liver radiotherapy. *Int. J. Radiat. Oncol. Biol. Phys.* **51**:1410–21; 2001.

76. Starkschall, G.; Himanshu, S.; Keall, P.J.; *et al.* Displacement-based binning of 4-D CT image data sets. XIVth International Conference on the Use of Computers in Radiation Therapy, 2004, pp 53.
77. Thorndyke, B.; Schreibmann, E.; Li, T.; *et al.* A comparison of amplitude- and phase-based 4D CT. *Med. Phys.* **32**:1919; 2005.
78. Lu, W.; Parikh, P.; Bradley, J.; *et al.* A comparison between amplitude sorting and phase sorting using external respiratory measurements for 4D CT. *Med. Phys.* **32**:2094; 2005.
79. Schreibmann, E.; Chen, G.T.Y.; Xing, L. Image interpolation in 4D CT using a BSpline deformable registration model. *Int. J. Radiat. Oncol. Biol. Phys.* 2006; In press.
80. Schweikard, A.; Shiomi, H.; Fisseler, J.; *et al.* Fiducial-less respiration tracking in radiosurgery. *LNCS* **3217**:992–99; 2004.
81. Vedam, S.S.; Kini, V.R.; Keall, P.J.; *et al.* Quantifying the predictability of diaphragm motion during respiration with a noninvasive external marker. *Med. Phys.* **30**:505; 2003.
82. Ford, E.C.; Mageras, G.S.; Yorke, E.; *et al.* Evaluation of respiratory movement during gated radiotherapy using film and electronic portal imaging. *Int. J. Radiat. Oncol. Biol. Phys.* **52**:522; 2002.
83. Yorke, E.; Rosenzweig, K.E.; Wagman, R.; *et al.* Interfractional anatomic variation in patients treated with respiration-gated radiotherapy. *J. Appl. Clin. Med. Phys.* **6**:19; 2005.
84. Klein, G.J.; Reutter, B.W.; Ho, M.H.; *et al.* Real-time system for respiratory-cardiac gating in positron tomography. *IEEE Trans. Nucl. Sci.* **45**:2139; 1998.
85. Nehmeh, S.A.; Erdi, Y.E.; Ling, C.C.; *et al.* Effect of respiratory gating on quantifying PET images of lung cancer. *J. Nucl. Med.* **43**:876–81; 2002.
86. Nehmeh, S.A.; Erdi, Y.E.; Rosenzweig, K.E.; *et al.* Reduction of respiratory motion artifacts in PET imaging of lung cancer by respiratory correlated dynamic PET: Methodology and comparison with respiratory gated PET. *J. Nucl. Med.* **44**:1644; 2003.
87. Cherry, S.R.; Dahlborn, M. Physics, Instrumentation and Scanners, PET: Molecular Imaging and its Biological Applications. New York: Springer-Verlag; 2004.
88. Sarikaya, I.; Yeung, H.W.; Erdi, Y.; *et al.* Respiratory artefact causing malpositioning of liver dome lesion in right lower lung. *Clin. Nucl. Med.* **28**:943; 2003.
89. Erdi, Y.E.; Nehmeh, S.A.; Pan, T.; *et al.* The CT motion quantitation of lung lesions and its impact on PET-measured SUVs. *J. Nucl. Med.* **45**:1287; 2004.
90. Pevsner, A.; Nehmeh, S.A.; Humm, J.L.; *et al.* Effect of motion on tracer activity determination in CT attenuation corrected PET images: A lung phantom study. *Med. Phys.* **32**:2358; 2005.
91. Nehmeh, S.A.; Erdi, Y.E.; Pan, T.; *et al.* Quantitation of respiratory motion during 4D-PET/CT acquisition. *Med. Phys.* **31**:1333–8; 2004.
92. Mattes, D.; Haynor, R.D.; Vesselle, H.; *et al.* PET-CT image registration in the chest using free-form deformations. *IEEE Trans. Med. Imaging* **22**:120–28; 2003.
93. Schreibmann, E.; Xing, L. Narrow band deformable registration of prostate MRI/MRSI and CT studies. *Int. J. Radiat. Oncol. Biol. Phys.* **62**:595–605; 2005.
94. Li, T.; Thorndyke, B.; Schreibmann, E.; *et al.* Model-based image reconstruction for four-dimensional PET. *Med. Phys.* Submitted.
95. Trofimov, A.; Rietzel, E.; Lu, H.M.; *et al.* Temporo-spatial IMRT optimization: Concepts, implementation and initial results. *Phys. Med. Biol.* **50**:2779–98; 2005.
96. Armbruster, B.; Yang, Y.; Schreibmann, E.; *et al.* Inverse planning for time-resolved aperture modulated radiation therapy. Annual ASTRO Meeting, Denver, CO, 2005.
97. Bortfeld, T.; Jiang, S.B.; Rietzel, E. Effects of motion on the total dose distribution. *Semin. Radiat. Oncol.* **14**:41–51; 2004.
98. Jiang, S.B.; Pope, C.; Al Jarrah, K.M.; *et al.* An experimental investigation on intra-fractional organ motion effects in lung IMRT treatments. *Phys. Med. Biol.* **48**:1773–84; 2003.
99. Rietzel, E.; Chen, G.T.; Choi, N.C.; *et al.* 4D image-based treatment planning: Target volume segmentation and dose calculation in the presence of respiratory motion. *Int. J. Radiat. Oncol. Biol. Phys.* **61**:1535; 2005.
100. Keall, P.J.; Joshi, S.; Vedam, S.S.; *et al.* Four-dimensional radiotherapy planning for DMLC-based respiratory motion tracking. *Med. Phys.* **32**:942–51; 2005.
101. Webb, S. The effect on IMRT conformality of elastic tissue movement and a practical suggestion for movement compensation via the modified dynamic multileaf collimator technique. *Phys. Med. Biol.* **50**:1163–90; 2005.
102. Yang, Y.; Xing, L. Clinical knowledge-based inverse treatment planning. *Phys. Med. Biol.* **49**:5101–17; 2004.
103. Yang, Y.; Xing, L. Inverse treatment planning with adaptively evolving voxel-dependent penalty scheme. *Med. Phys.* **31**:2839–44; 2004.
104. Wu, Q.; Mohan, R.; Niemierko, A.; *et al.* Optimization of intensity-modulated radiotherapy plans based on the equivalent uniform dose. *Int. J. Radiat. Oncol. Biol. Phys.* **52**:224–35; 2002.
105. Olivera, G.H.; Shepard, D.M.; Reckwerdt, P.J.; *et al.* Maximum likelihood as a common computational framework in tomotherapy. *Phys. Med. Biol.* **43**:3277–94; 1998.
106. Lee, E.K.; Gallagher, R.J.; Silvern, D.; *et al.* Treatment planning for brachytherapy: An integer programming model, two computational approaches and experiments with permanent prostate implant planning. *Phys. Med. Biol.* **44**:145–65; 1999.
107. Shepard, D.M.; Earl, M.A.; Li, X.A.; *et al.* Direct aperture optimization: A turnkey solution for step-and-shoot IMRT. *Med. Phys.* **29**:1007–18; 2002.
108. Bednarz, G.; Michalski, D.; Houser, C.; *et al.* The use of mixed-integer programming for inverse treatment planning with predefined field segments. *Phys. Med. Biol.* **47**:2235–45; 2002.
109. Cotrutz, C.; Xing, L. Segment-based dose optimization using a genetic algorithm. *Phys. Med. Biol.* **48**:2987–98; 2003.
110. Murphy, M.J.; Adler, Jr., J.R.; Bodduluri, M.; *et al.* Image-guided radiosurgery for the spine and pancreas. *Comput. Aided Surg.* **5**:278–88; 2000.
111. Murphy, M.J. Tracking moving organs in real time. *Semin. Radiat. Oncol.* **14**:91–100; 2004.
112. Sharp, G.C.; Jiang, S.B.; Shimizu, S.; *et al.* Prediction of respiratory tumour motion for real-time image-guided radiotherapy. *Phys. Med. Biol.* **49**:425–40; 2004.
113. Berbeco, R.I.; Neicu, T.; Rietzel, E.; *et al.* A technique for respiratory-gated radiotherapy treatment verification with an EPID in cine mode. *Phys. Med. Biol.* **50**:3669–79; 2005.
114. Herman, M.G.; Balter, J.M.; Jaffray, D.A.; *et al.* Clinical use of electronic portal imaging: Report of AAPM Radiation Therapy Committee Task Group 58. *Med. Phys.* **28**:712–37; 2001.
115. Yang, Y.; Xing, L. Quantitative measurement of MLC leaf displacements using an electronic portal image device. *Phys. Med. Biol.* **49**:1521–33; 2004.
116. Yan, D.; Wong, J.; Vicini, F.; *et al.* Adaptive modification of treatment planning to minimize the deleterious effects of treatment setup errors. *Int. J. Radiat. Oncol. Biol. Phys.* **38**:197–206; 1997.
117. Court, L.E.; Dong, L.; Taylor, N.; *et al.* Evaluation of a contour-alignment technique for CT-guided prostate radiotherapy: An intra- and interobserver study. *Int. J. Radiat. Oncol. Biol. Phys.* **59**:412–8; 2004.
118. Orton, N.P.; Tome, W.A. The impact of daily shifts on prostate IMRT dose distributions. *Med. Phys.* **31**:2845–8; 2004.
119. Fuss, M.; Salter, B.J.; Cavanaugh, S.X.; *et al.* Daily ultrasound-based image-guided targeting for radiotherapy of upper abdominal malignancies. *Int. J. Radiat. Oncol. Biol. Phys.* **59**:1245–56; 2004.
120. Saw, C.B.; Ayyangar, K.M.; Zhen, W.; *et al.* Clinical implementation of intensity-modulated radiation therapy. *Med. Dosim.* **27**:161–9; 2002.
121. Schweikard, A.; Glosser, G.; Bodduluri, M.; *et al.* Robotic motion compensation for respiratory movement during radiosurgery. *Comput. Aided Surg.* **5**:263–77; 2000.
122. Mackie, T.R.; Balog, J.; Ruchala, K.; *et al.* Tomotherapy. *Semin. Radiat. Oncol.* **9**:108–17; 1999.
123. Meeks, S.; Harmon, J.J.; Langen, K.; *et al.* Performance characterization of megavoltage computed tomography imaging on a helical tomotherapy unit. *Med. Phys.* **32**:2673–81; 2005.
124. Jaffray, D.A.; Siewersen, J.H.; Wong, J.W.; *et al.* Flat-panel cone-beam computed tomography for image-guided radiation therapy. *Int. J. Radiat. Oncol. Biol. Phys.* **53**:1337–49; 2002.
125. Pouliot, J.; Bani-Hashemi, A.; Chen, J.; *et al.* Low-dose megavoltage cone-beam CT for radiation therapy. *Int. J. Radiat. Oncol. Biol. Phys.* **61**:552–60; 2005.

126. Oldham, M.; Letourneau, D.; Watt, L.; et al. Cone-beam-CT guided radiation therapy: A model for on-line application. *Radiat. Oncol.* **75**:271–8; 2005.
127. Li, T.; Yang, Y.; Schreibmann, E.; et al. A new cone-beam CT reposition technique through deformable registration. Annual Meeting of American Association of Physicists in Medicine, Seattle, WA, 2005.
128. Mohan, R.; Zhang, X.; Wang, H.; et al. Use of deformed intensity distributions for on-line modification of image-guided IMRT to account for interfractional anatomic changes. *Int. J. Radiat. Oncol. Biol. Phys.* **61**:1258–66; 2005.
129. Yang, Y.; Li, T.; Schreibmann, E.; et al. Is cone beam CT suitable for dose verification? Annual Meeting of American Association of Physicists in Medicine, Seattle, WA, 2005.
130. Yang, Y.; Schreibmann, E.; Xing, L.; et al. Dosimetric evaluation of kV cone-beam CT (CBCT)-based dose calculation. *Phys. Med. Biol.* 2006; Submitted.
131. Wang, G.; Vannier, M.W. Preliminary study on helical CT algorithms for patient motion estimation and compensation. *IEEE Trans. Med. Imaging* **14**:205–11; 1995.
132. Willis, N.P.; Bresler, Y. Optimal scan for time-varying tomography. I: Theoretical analysis and fundamental limitations *IEEE Trans. Med. Imaging* **14**:642–53; 1995.
133. Atalar, E.; Onural, L. A respiratory motion artifact reduction method in magnetic resonance imaging of the chest. *IEEE Trans. Med. Imaging* **10**:11–24; 1991.
134. Crawford, C.; King, K.F.; Ritchie, C.J.; et al. Respiratory compensation in projection imaging using a magnification and displacement model. *IEEE Trans. Med. Imaging* **15**:327–32; 1996.
135. Li, T.; Schreibmann, E.; Yang, Y.; Xing, L. Respiratory motion correction with deformation field for improved on-board imaging with cone-beam CT. *Phys. Med. Biol.* **51**:253–67; 2006.
136. Sonke, J.J.; Zijp, L.; Remeijer, P.; et al. Respiratory correlated cone beam CT. *Med. Phys.* **32**:1176–86; 2005.
137. Li, T.; Xing, L.; Munro, P.; et al. 4D cone beam CT using an on-board imager. *Med. Phys.* 2006; Submitted.
138. Zeng, R.; Fessler, J.A.; Balter, J.M. Respiratory motion estimation from slowly rotating x-ray projections: Theory and simulation. *Med. Phys.* **32**:984–91; 2005.
139. Court, L.E.; Dong, L. Automatic registration of the prostate for computed-tomography-guided radiotherapy. *Med. Phys.* **30**:2750–7; 2003.
140. Bharath, A.; Hirose, M.; Hata, N.; et al. Evaluation of three-dimensional finite element-based deformable registration of pre- and intraoperative prostate imaging. *Med. Phys.* **28**:2551–60; 2001.
141. Brock, K.K.; McShan, D.L.; Balter, J.M. A comparison of computer-controlled versus manual on-line patient setup adjustment. *J. Appl. Clin. Med. Phys.* **3**:241–7; 2002.
142. Bookstein, F.L. Thin plate splines and the decomposition of deformations. *IEEE Trans. Pattern Anal. Mach. Intell.* **11**:567–85; 1989.
143. Lian, J.; Xing, L.; Hunjan, S.; et al. Mapping of the prostate in endorectal coil-based MRI/MRSI and CT: A deformable registration and validation study. *Med. Phys.* **31**:3087–94; 2004.
144. Guerrero, T.; Zhang, G.; Huang, T.C.; et al. Intrathoracic tumour motion estimation from CT imaging using the 3D optical flow method. *Phys. Med. Biol.* **49**:4147–61; 2004.
145. Mageras, G.S.; Pevsner, A.; Yorke, E.D.; et al. Measurement of lung tumor motion using respiration-correlated CT. *Int. J. Radiat. Oncol. Biol. Phys.* **60**:933–41; 2004.
146. Schwartz, D.L.; Ford, E.C.; Rajendran, J.; et al. FDG-PET/CT-guided intensity modulated head and neck radiotherapy: A pilot investigation. *Head Neck* **27**:478–87; 2005.
147. Fei, B.; Wheaton, A.; Lee, Z.; et al. Automatic MR volume registration and its evaluation for the pelvis and prostate. *Phys. Med. Biol.* **2002**:823–38; 2002.
148. Coselman, M.M.; Balter, J.M.; McShan, D.L.; et al. Mutual information based CT registration of the lung at exhale and inhale breathing states using thin-plate splines. *Med. Phys.* **31**:2942–8; 2004.
149. Schreibmann, E.; Xing, L. Image registration with auto-mapped control volumes. *Med. Phys.* Submitted.
150. Koong, A.C.; Le, Q.T.; Ho, A.; et al. Phase I study of stereotactic radiosurgery in patients with locally advanced pancreatic cancer. *Int. J. Radiat. Oncol. Biol. Phys.* **58**:1017–21; 2004.
151. Loo, B.; Thorndyke, B.; Maxim, P.; et al. Determining margin for target deformation and rotation in respiratory motion-tracked stereotactic radiosurgery of pancreatic cancer. Annual ASTRO Meeting, Denver, CO, 2005.
152. Koong, A.C.; Le, Q.T.; Ho, A.; et al. Phase II study of stereotactic radiosurgery in patients with locally advanced pancreatic cancer. *Int. J. Radiat. Oncol. Biol. Phys.* In press.

Evaluation of Onboard kV Cone Beam CT (CBCT)- Based Dose Calculation[†]

Yong Yang, Eduard Schreibmann, Tianfang Li, Chuang Wang,
and Lei Xing^{a)}

Department of Radiation Oncology, Stanford University School of Medicine, 875 Blake
Wilbur Drive, Stanford, CA 94305-5847

^{a)} Author to whom correspondence should be addressed:
Department of Radiation Oncology
Stanford University School of Medicine,
Clinical Cancer Center
875 Blake Wilbur Drive, Rm G-204
Stanford, CA 94305-5847
Telephone: (650) 498-7896
Fax: (650) 498-4015
Email: lei@reyes.stanford.edu

Submitted to: *Physics in Medicine and Biology*

[†] Part of this work was presented in 2006 Annual Meeting of American Association of
Physicists in Medicine.

Abstract: On-board CBCT images are used to generate patient geometric model to assist patient setup. The image data can also, potentially, be used for dose reconstruction in combination with the fluence maps from treatment plan. Here we evaluate the achievable accuracy in using a kV CBCT for dose calculation. Relative electron density as a function of HU was obtained for both planning CT (pCT) and CBCT using a Catphan-600 calibration phantom. The CBCT calibration stability was monitored weekly for 8 consecutive weeks. A clinical treatment planning system was employed for pCT- and CBCT-based dose calculations and subsequent comparisons. Phantom and patient studies were carried out. In the former study, both Catphan-600 and pelvic phantoms were employed to evaluate the dosimetric performance of the full-fan and half-fan scanning modes. To evaluate the dosimetric influence of motion artifacts commonly seen in CBCT images, the Catphan-600 phantom was scanned with and without cyclic motion using the pCT and CBCT scanners. The doses computed based on the four sets of CT images (pCT and CBCT with/without motion) were compared quantitatively. The patient studies included a lung case and three prostate cases. The lung case was employed to further assess the adverse effect of intra-scan organ motion. Unlike the phantom study, the pCT of a patient is generally acquired at the time of simulation and the anatomy may be different from that of CBCT acquired at the time of treatment delivery because of organ deformation. To tackle the problem, we introduced a set of modified CBCT images (mCBCT) for each patient, which possesses the geometric information of the CBCT but the electronic density distribution mapped from the pCT with the help of a BSpline deformable image registration software. In the patient study, the dose computed with the mCBCT was used as a surrogate of the “ground truth”.

We found that the CBCT electron density calibration curve differs moderately from that of pCT. No significant fluctuation was observed in the calibration over the period of 8 weeks. For the static phantom, the doses computed based on pCT and CBCT agreed to within 1%. Notable difference in CBCT- and pCT-based dose distributions was found for the motion phantom due to the motion artifacts appeared in the CBCT images (the maximum discrepancy was found to be ~3.0% in the high dose region). The motion artifacts-induced dosimetric inaccuracy was also observed in the lung patient study. For the prostate cases, the mCBCT- and CBCT-based dose calculations yielded very close results (<2%). Coupled with the phantom data, it is concluded that the CBCT can be employed directly for dose calculation for a disease site such as prostate, where there is little motion artifact. In the prostate case study, we also noted a large discrepancy between

the original treatment plan and the CBCT (or mCBCT)-based calculation, suggesting the importance of inter-fractional organ movement and the need for adaptive therapy to compensate for the anatomical changes in the future.

Key word: CBCT, IGRT, Dose verification, Deformable registration, IMRT

I. Introduction

Modern radiation therapy techniques, such as 3D conformal radiotherapy (3DCRT) and intensity-modulated radiation therapy (IMRT), provide unprecedented means for producing exquisitely shaped radiation doses that closely conform to the tumor dimensions while sparing sensitive structures. As a result of greatly enhanced dose conformality, more accurate beam targeting becomes an urgent issue in radiation therapy. In current practice, uncertainties exist in tumor target localization due to intra- and inter-organ motions during the course of radiation treatment. As thus, large safety margins around the tumor targets and sensitive structures are introduced to cope with the otherwise insoluble patient localization problem. The use of non-optimal margins compromises the patient care and adversely affects the treatment outcome (Cai et al., 2006; Coolens et al., 2006; Hector et al., 2000; Hugo et al., 2003; Litzenberg et al., 2006; Ruan et al., 2006; Webb, 2006). The need to improve targeting in high precision radiation therapy has recently spurred a flood of research activities in image-guided radiation therapy (IGRT) (Xing et al., 2006).

CBCT based upon flat-panel technology integrated with a medical linear accelerator has recently become available from Linac vendors for therapy guidance. The volumetric images may be used to verify and correct the planning patient setup in the linac coordinates by comparing with the patient position defined in treatment plan. Both kV and MV beams have been utilized for this application (Chang et al., 2006; Dietrich et al., 2006; Hawkins et al., 2006; Jaffray et al., 2002; Langen et al., 2005; Sorcini and Tilikidis, 2006; Thilmann et al., 2006; Yin et al., 2005). The former typically consists of a kV-source and flat-panel combination mounted on the drum of a medical accelerator, with the kV imaging axis orthogonal to that of MV therapy beam. In addition to guide the patient setup process, CBCT data acquired prior to the treatment can, in principle, be used to recalculate or verify the treatment plan based on the patient anatomy of the treatment day. The recalculation starts with the intended fluence maps from the patient's treatment plan, whereas the verification is done by using the fluence maps measured at the exiting sides of the incident beams. If CBCT-based dose calculation is accurate enough (say, with an accuracy within 1~2%), this will provide a valuable option for us to

predict/assess the patient dose routinely. In reality, because of the presence of organ movement/deformation, it is conceivable that the dose distributions delivered to the patient are usually different from fraction to fraction. It is paramount to be able to monitor the actual patient dose for each fraction as well as the accumulated doses to the target and sensitive structures while the fractional treatment proceeding. This will not only give physician more confidence about the treatment but may, in future, afford us an effective means to adaptively modify the patient's treatment plan during the course of a radiation therapy based on the dose that has already been delivered.

Different from conventional CT, kV CBCT covers a much larger field of view (FOV) in the longitudinal direction, and scatter poses a more severe problem in the resultant image. In addition, the gantry rotation speed is limited to ~1 min. by IEC regulation, which makes the CBCT more prone of motion artifacts. The deteriorated image quality raises serious concern about the dosimetric reliability of CBCT-based dose calculation. In this work we evaluate the dosimetric accuracy of kV CBCT-based dose calculation.

II. Method and Materials

A. Data acquisition

The on-board imager (OBI) integrated in a Trilogy™ medical linear accelerator (Varian Medical Systems, Palo Alto, CA) is used in this work to acquire CBCT images. The kV OBI system is capable of obtaining low-dose, high-resolution radiography, fluoroscopy and CBCT. The system is mounted on the treatment machine via robotically controlled arms, which operate along three axes of motion. A 150 kV X-ray tube with maximum 32 ms pulse length for continuous irradiation and maximum 320 ms pulse length for single pulse is designed for generating high-resolution images from a moving gantry. The spot of the tube is located at 90° to the MV source and 100cm from the radiation axis of the accelerator. A 39.7cm X 29.8 cm amorphous silicon flat-panel X-ray image detector (Varian PortalVision™ aS1000) mounted opposite the kV tube is used to acquire digital images with a pixel matrix of 2048 X 1536. Using the OBI system, the CBCT data can be

acquired in two modes: full fan mode and half fan mode. In the full fan mode, the beam central axis passes through the detector center and a full projection of the scanned patient is acquired for each acquirement angle. Total 675 projections are taken during the whole 364° gantry rotation with a maximum field of view (FOV) about 25 cm in diameter and 17cm in length. The data acquisition time is about 60 second and the reconstruction time for 340 slices of 512X512 CBCT images with a voxel size of 0.5mm³ is also about 1 minute on a PC. The half fan mode is designed to obtain larger FOV. In this mode, the detector is shifted laterally to take only half of the projection of the scanned patient for each acquirement angle. Total about 965 projections are taken during the 364° gantry rotation and a FOV of 50 cm in the axial plane and 15cm in the longitudinal direction can be achieved. The data acquisition and reconstruction time for 512X512 CBCT images with a voxel size of 0.95mm³ using this mode is about double compared with the full fan mode. The averaged dose for a head and neck CBCT scan is about 2 cGy, and 3cGy for an abdominal scan.

B. Calibration of relative electron density

To use CT or CBCT for radiation dose calculation, it is required to relate the Hounsfield Unit (HU) of the scanner with the actual electron density. A CT-phantom, Catphan-600 module CTP404 (Phantom Laboratory, NY), was used for the calibration of pCT (GE Discovery-ST PET/CT scanner, Milwaukee, WI) and CBCT. The gantry rotation speed of the 16-slice Discovery-ST scanner is 0.5sec/rotation. The CTP404 has a diameter of 150 mm (the longitudinal dimension of the phantom is 16cm) and contains 17 different sizes of inserts with seven different tissue substitute materials, air, PMP, LDPE, Polystyrene, Acrylic, Delrin and Teflon, respectively. Their relative electron densities ranged from 0 to 1.866. A cross section of the phantom is shown in figure 1. The calibration of a CT scanner involves acquiring CT images, obtaining average HUs for each inserting materials, and plot the HU as a function of the relative electron density. For CBCT calibration, the only difference from the conventional CT is that it is necessary to calibrate separately for full and half fan modes because the beam geometry and characteristics of the two types of scanning modes are different.

In order to test the stability of the CBCT calibration curve with time, the phantom was repeatedly scanned every week for two months. The obtained HU vs relative electron density curves were compared to assess the HU fluctuations with time.

C. Phantom study

CT and CBCT images of the Catphan-600 phantom were acquired using the procedure outlined in Sec. II.A. The phantom was placed on a platform that can be set to one-dimensional cyclic motion with a number of speeds. The axis of the cylindrically shaped phantom, along which the phantom moves cyclically, was angled from the central axis of the CBCT gantry rotation by about 15° in order to study the motion influence on CT/CBCT. The movement of the phantom produces motion artifacts in the images and allows us to evaluate the performance of CBCT-based dose calculation in the presence of organ motion. The full fan mode was used to scan the phantom. CT and CBCT images of the phantom were acquired with and without motion. In the former case, the peak-to-peak amplitude of the motion was 0.5 cm in the left-right direction and the period was 4s. In addition, different sizes of homogeneity cylindrical phantoms, with a diameter of 10.8, 16, and 26.6 cm, respectively, were scanned to evaluate the scatter influence on image quality.

To quantify the difference in the image quality of the CT and CBCT images, we first analyzed the HU distribution for all the acquired images. The influence of phantom motion and scatter radiation on the HU distribution was investigated. The CT and CBCT images were imported to a Varian Eclipse treatment planning system for dosimetric comparison study. For planning and evaluation purpose, a hypothetical spherical target with a diameter of 5cm was created at the center of the phantom and a single 5 X 5cm² 6MV photon beam was used to irradiate the target. A simple beam configuration was used here because, in this way, the results are more intuitively interpretable. The pencil beam convolution dose calculation algorithm implemented in Varian Eclipse treatment planning system was adopted for dose calculation. The resultant isodose curves, dose profiles and DVHs were compared.

A detailed study using a CIRS Model 002PRA pelvic IMRT phantom (CIRS Tissue Simulation & Phantom Technology, Norfolk, Virginia) was also performed to

validate the on-board CBCT-based dose calculation. The specifications of the phantom can be found from http://www.cirsinc.com/002pra_rad.html. After CT and CBCT images were acquired, a five-field IMRT plan was generated with hypothetical target and sensitive structure at the center of the phantom. The dose distributions and DVHs obtained using the two sets of images were then compared.

D. Patient study

Three prostate and a lung patients were selected for the evaluation study of CBCT-based dose calculation. For the prostate cases, the targets included the PTV, consisting of the prostate gland with a margin of 6mm, and the seminal vesicles (SVs). The critical structures were rectum, bladder and femoral heads. IMRT plans using five 15MV photon beams with gantry angles of 35°, 110°, 180°, 250°, and 325° (in IEC convention) were adopted for the three cases. All the plans were normalized to deliver a prescription dose of 78Gy to 99% the prostate PTV and no less than 50Gy to the 98% of SVs in 39 fractions.

The patient was setup by using kV on-board imager (orthogonal projection images) under the guidance of implanted fiducials. After the kV AP/LAT images were acquired, the patient is shifted in such a way that the implanted gold seeds match with that on the DRRs. After the patients were setup using the current clinical procedure, CBCT images of the patients were acquired. The CBCT images were transferred to an Eclipse treatment planning system (Varian Medical Systems, Palo Alto, CA). For each case, the IMRT planning parameters generated for the patient's treatment, including beam configuration, MU settings, and fluence maps, were employed to recalculate the dose based on the CBCT data. The CT and CBCT-based treatment plans were then compared.

For the lung cancer case, the PTV consisted of the CTV with a margin of 10mm. The critical structures involved were the right lung and the spinal cord. A conventional 3D conformal plan with three 15MV photon beams (45°, 180° and 288° in IEC convention) was generated for this case. In this plan, the field shape of each beam was determined by conforming the PTV projection in the corresponding beam direction. The plan was normalized to deliver a prescription dose of 70Gy to 100% of the target volume

in 35 fractions. The CBCT images of the patients were obtained using the half fan mod.

E. Deformable electron density mapping (DEDM)

Unlike the phantom study, the pCT of a patient is generally acquired at the time of simulation and the anatomy may be different from that of CBCT acquired at the time of treatment delivery because of organ deformation. This makes it difficult to assess the accuracy of CBCT-based dose reconstruction. To tackle the problem, for each patient, we introduced a set of modified CBCT images (mCBCT) and the dose computed based on the mCBCT was used as a surrogate of the “ground truth” in the evaluation of CBCT-based dose calculation. The mCBCT possesses the geometric information of the CBCT yet the electronic density distribution mapped from the pCT. The correspondence between the pCT and CBCT was accomplished by using an in-house BSpline deformable image registration software. The electron density mapping process is described as follows.

A free form spline (BSpline) deformable model (Li et al., 2005; Li et al., 2006; Paquin et al., 2006; Schreibmann et al., 2006) was employed to register the pCT and CBCT. Briefly, a lattice of user-defined nodes is overlaid on the image. Each node contains a deformation vector, whose components are determined by optimizing a metric function that characterizes the goodness of the registration. In this work, a voxel-based normal cross correlation (NCC) metric was used and the node deformations were determined using the gradient-based optimization algorithm. The deformation at any point of the image is calculated by spline interpolation of closest nodes values. After registration, the HU in each voxel in pCT was mapped to the corresponding point in the CBCT to produce the mCBCT images.

III. Results

A. Calibration of CT and CBCT

The relationship between kV HU distribution of CBCT and relative electron dosimetry was established by using a Catphan-600 CT phantom following the procedure described in Sec. II.B. The calibration curves for pCT, half fan and full fan CBCT modes are shown

in figure 2a. Figure 2b shows the calibration data obtained with an interval of 1 week during a period of two months for full fan CBCT. No significant variations were found in the calibration, which is similar to what have been observed for MV CBCT. The stability of the kV CBCT and electron density calibration is a good indicator of the HU number integrity and the overall performance of the CBCT system.

B. Phantom study

Figures 1a to 1d show the same transverse slices of the CT and CBCT images of the Catphan-600 phantom with and without motion. The first two panels are the CT and CBCT images of the phantom in the absence of motion, and the second two show the same with the phantom motion “switched on”. It is seen that the quality of CBCT images is worse than that of the conventional pCT, especially in the presence of motion. The HU profiles of the four images along the two orthogonal lines (lines A-A and B-B as marked in figure 1) are plotted in figure 3. It is found that the HU profiles of the pCT and CBCT normally agree to within 10% in the static situation. On the other hand, when the motion is “switched on”, CBCT shows a much greater level of artifacts (figure 1d) and the HU difference between the conventional CT and CBCT is aggravated, with the maximum difference reaching several hundred HUs.

Because of the cone beam geometry, the influence of scatter radiation in CBCT is more severe as compared to that of a fan beam geometry. In general, X-ray scatter reduces image contrast, increases image noise and may introduce reconstruction error into CBCT. Figure 4 plots the HU profiles along a central axis of three different sized homogeneous cylindrical water phantoms. As expected, the fluctuation range of HU value increases with the phantom size, indicating the increased influence of scatter radiation.

Figures 5, 6 and 7 present the dosimetric results calculated using a single 6 MV 5 X 5cm² photon beam. Figures 5a to 5d depict the dose distributions in a transverse slice calculated based on the four sets of images given in figure 1. Figures 6a and 6b compare the dose profiles along the two orthogonal lines (lines A-A and B-B in figure 1), and figure 7 compares the DVHs of the target for the four different situations. From these results we see that the dose calculated using pCT agrees with that of CBCT-based

calculation to within 1.0%, indicating that it is acceptable to use kV CBCT for dose calculation if no organ motion presents. However, when phantom motion is involved, the motion-induced artifacts significantly influence the HU distribution and thus the accuracy of CBCT-based dose calculation. For this simple phantom case, we find that the discrepancy between the pCT- and CBCT-based calculations is about 3%, which is clinically significant. Calculations using a few other field sizes ranging from 5 to 20 cm showed similar level of agreement in the case of static phantom and the results are not presented here due to space limitation. The motion artifacts existing in current CBCT limit the direct use of CBCT for dose calculation when intra-scanning organ motion is not negligible.

Figures 8 compares the IMRT dose calculations computed based on pCT and CBCT in the CIRS pelvic phantom. The DVHs of the target and sensitive structure are also presented in figure 8. Once again, it is seen that the dose calculated using pCT agrees with that of CBCT-based calculation to within 1.0%, indicating that it is acceptable to use kV CBCT for dose calculation for a pelvic sized phantom.

C. Patient study

Figures 9a to 9c show the same transverse slices of the pCT, CBCT, and checkerboard image resulting from the deformable registration of the two sets of images for one of the prostate cases. The mCBCT obtained by mapping the HUs from the pCT to CBCT is shown in figure 9d. Our previous studies have indicated that an accuracy better than 3mm is achievable by using the BSpline deformable model even in the presence of large deformations (Balter et al., 2005; Cai et al., 2006; Dietrich et al., 2006; Kamath et al., 2006). As can be seen from the checkerboard overlay, the registration between CT and CBCT is excellent. Figure 10 shows the isodose distributions for the three calculations based on pCT, CBCT, and mCBCT for the same case. A comparison of DVHs of PTV, prostate, SVs, bladder and rectum is presented in bottom of figure 10. Figures 11 and 12 present the DVHs of the targets and sensitive structures for the second and third cases. While there is significant dosimetric discrepancy between the pCT- and mCBCT-based plans, the results obtained using the CBCT or mCBCT is similar. The differences

between the CBCT- and mCBCT-based calculations are generally less than 2%, suggesting that it is suitable to directly use CBCT for dose recalculation.

For all three prostate cases, we found that the CBCT-reconstructed prostate dose agrees with the planned one to within ~3%. In Figure 10, the maximum prostate doses in the calculations based on planning CT, CBCT and modified CBCT are 105.3%, 107.2% and 106.4%, respectively. The maximum doses obtained with planning CT or CBCT differ from that obtained using modified CT is only 1%. In the absence of registration error, the dose difference between the planning CT and modified CBCT is due to daily anatomical change while the difference between the CBCT and modified CBCT is attributed to the HU value change. We found the dosimetric difference in the PTV is quite significant, which could be greater than 6~10%. Similar observation was made in studies using daily CT on-rail (Court et al., 2005; Paskalev et al., 2004). For the rectum and bladder, the discrepancies between planned and reconstructed doses could be greater than 8%. We attribute the discrepancy between the planned and CBCT-reconstructed dose distribution to the inter-fractional organ movements.

Figures 13a to 13d show the same transverse slice from the planning CT, CBCT, checkerboard overlay of the planning CT and the mCBCT for the lung cancer case. Figure 14 compares the isodose distributions and the DVHs of the target and sensitive structures calculated based on the three sets of images for the 3D conformal treatment. In this case, the dosimetric discrepancy between the CBCT- and mCBCT-based calculations is larger than that in the prostate cases, especially for the right lung and PTV. The maximum dose difference is about 5%. However, the discrepancy between the results obtained using planning CT and mCBCT becomes much less. This study exemplifies that, unlike the prostate cases, the dosimetric inaccuracy arises from the inferior image quality of CBCT in the dose verification calculation. The motion artifacts not only make it difficult to see the extent of the tumor, but also limit the direct use of CBCT for dose calculation in clinical practice.

IV. Discussion

CBCT volumetric imaging integrated with a medical linear accelerator opens new avenues for improving current radiation oncology practice. CBCT has two important applications: patient setup and dose reconstruction/verification. By imaging the patient routinely during a course of radiation therapy, the accuracy of the patient setup can potentially be improved. Furthermore, the CBCT provides a pre-treatment patient model upon which the dose calculation can be performed using intended fluence maps from the planning system or other means. Both applications rely on quality of the volumetric images. In reality, the design of current on-board CBCT is far from optimal. Its quality is adversely influenced by many factors, such as scatter, beam hardening, and intra-scanning organ motion. This raises the question whether current CBCT images can be used directly for radiation dose calculation. In this work, the feasibility and accuracy of using a commercial kV CBCT to calculate dose are investigated with a CT phantom and three clinical prostate cases.

Dose distributions computed based on CBCT or mCBCT represent the dose to be delivered to the patient because the CBCT is usually acquired prior to the patient's treatment after repositioning under the standard patient setup procedure. In general, the difference between the planned and CBCT-reconstructed dose distributions arises from two factors: (i) patient positioning error and organ deformation/displacement; and (ii) relative electron density variation in CBCT images. The small discrepancy between the doses computed using CBCT and mCBCT (see last section) suggests that, in the prostate cases, the second factor is small and it is acceptable to directly use CBCT for dose calculation. Our static phantom studies also support the notion. However, the dosimetric error caused by the inter-scanning organ motion/deformation is not insignificant, as suggested by the large discrepancies between the original plan and the dose distribution reconstructed based CBCT or mCBCT. For example, in the first case, the SVs receive ~50% lower dose as compared to the planned delivery, whereas the second shows ~10% higher dose as compared to the planned one. In the third case, a clinically unacceptable SV underdosing is observed. The significant difference in SV DVHs of the original and to-be-delivered plans is caused by the relatively small volume and mobile nature of the

SV. For the rectum and bladder, the discrepancies between planned and reconstructed doses are less yet clinically significant in all three cases.

While CBCT-based dose agree with mCBCT-based calculation to within ~2% in most situations and the CBCT images can be used directly for dosimetric validation for a disease site such as the prostate, it is not recommended to replace the conventional planning CT by CBCT for the purpose of radiation treatment planning. In addition to the inferior image quality, which may affect the physician's ability to delineate the tumor and sensitive structures, the limited FOV of CBCT (especially in the longitudinal direction) often renders the resultant images non-useable for treatment planning.

A few groups are working on deformable model based segmentation and patient setup procedures (Cai et al., 2006). When deformable model is used, one can go beyond simply aligning the 3D bony structures to achieve a registration based on the matching of soft-tissue organ(s). There are multiple ways to set up the patient depending on which structure to align during the patient positioning. Generally, the degree of freedom available for the patient setup is much less than the dimensionality of the patient deformation. Thus a deformable registration based patient setup is not likely to be the ultimate solution to volumetric image-guided radiation therapy, even though it improves the current bony-structure-based patient alignment by partially taking into account organ deformation. A more attractive solution to accommodate various factors mentioned above is to re-optimize (or tweak) the IMRT plan based on the patient's setup CBCT. In order to fully utilize the volumetric imaging data, a new paradigm with seamlessly integrated simulation, planning, verification, and delivery procedure is needed. Until this is realized clinically, the volumetric image-based patient positioning is simply an expensive extension of the existing planar setup procedure.

In the DEDM technique described above, pre-assumption is that the accurate registration between pCT and CBCT is achievable. In reality, an excellent registration is attainable in the case of rigid bodies through a variety of means, such as fiducials, manual, or automated registration method. The validity and advantage of the approach in this situation can, therefore, be easily envisaged. In the phantom case presented earlier, it is perceivable that the CBCT-phantom dose calculation can be trivially improved if the pCT imaging data (figure 5b) are mapped to CBCT (figure 5a). The existence of organ

deformation in more general situations necessitates a deformable registration for electron density mapping. The BSpline method used in this work was employed for several IGRT related projects involving intra- and inter-modality deformable registrations and its performance has been demonstrated (Balter et al., 2005; Cai et al., 2006; Dietrich et al., 2006; Kamath et al., 2006). The studies carried out by our group and others (Allen et al., 2004; Brock et al., 2003; Coselmon et al., 2004; Xiong et al., 2006) have suggested that a spatial accuracy better than 3mm is readily achievable using various deformable image registration techniques. The accuracy can be further improved when more sophisticated methods are used (Kamath et al., 2006). To assess the influence of registration accuracy on the mCBCT-based dose reconstruction, for the first prostate case, we intentionally introduced a number of “wrong-assignments” of the electron density distribution when mapping the electron density from pCT to CBCT (the spatial inaccuracy of the “wrong-assignments” was restricted to within 3mm in the soft tissue regions in these digital experiments). The resultant dosimetric deviations from the presented mCBCT-based calculations were found negligible (all within 1.5%), which indicate that (i) the final dose distribution is not very sensitive to small regional variations in the registration; and (ii) mCBCT-reconstructed dose represents a good “surrogate” of the “ground truth”. The fact that the CBCT- and mCBCT-based dose calculations differ little also supports the above arguments. A more direct and perhaps the ultimate validation of the DEDM could be accomplished if the patient had undergone a CT on-rail scan (right before or after the CBCT acquisition) in addition to the pCT and CBCT. Before this type of data become available, we alert the readers that the DEDM approach remains heuristic in nature and one should bear its inherent limitation in mind when using it for dose calculation.

In the presence of intra-scan organ motion, our phantom study indicated that significant dosimetric errors could happen. The same conclusion can be drawn from the lung study. Recent developments of 4D CT (Allen et al., 2004; Berbeco et al., 2005; de Koste et al., 2003; Vedam et al., 2003; Wink et al., 2005) and 4D CBCT (Dietrich et al., 2006; Sonke et al., 2005; Li et al. 2006), in conjunction with the proposed DEDM method, may provide a valuable solution to the problem of CBCT-based dose calculation in the regions of thorax and upper abdomen. This is clearly a subject of future investigation.

It is worth to mention that previous studies indicate that bulk CT density correction can be quite effective for heterogeneity corrections for all treatment sites. For prostate case treated with high energy x-rays, the homogeneous dose calculation method will only introduce 2-4% error (depending on beam directions and relative weights). It may be acceptable to use homogeneous dose calculation method for certain cases, as suggested by a study by Chen et al (Chen et al., 2004). A detailed study along this line should be useful in the future.

V. Conclusion

On-board CBCT provides useful volumetric anatomy information for patient positioning verification. When used for dose verification calculation, a reliable relationship between the HU and relative electron density is needed. Our phantom and prostate patient studies have indicated that, in the absence of motion artifacts, the dosimetric accuracy of CBCT-based dose calculation is acceptable for the purpose of dosimetric check. Our motion phantom and lung patient studies showed that the dosimetric errors become more pronounced when intra-scanning organ motion is present. In this situation, a direct use of CBCT for dose calculation is not recommended. Before an effective CBCT image quality improving technique is in place, DEDM may be a useful interim solution for improved dose calculation.

Acknowledgement

We wish to acknowledge Drs. C. King, S. Hancock, T. Pawlicki, P. Maxim, and G. Luxton for useful discussions. The supports from the National Cancer Institute (1 R01 CA98523-01), Department of Defense (PC040282), and Varian Medical Systems are also gratefully acknowledged.

References

- Allen AM, Siracuse KM, Hayman JA, Balter JM. 2004. Evaluation of the influence of breathing on the movement and modeling of lung tumors. *Int J Radiat Oncol Biol Phys* 58:1251-1257.
- Balter JM, Wright JN, Newell LJ, Friemel B, Dimmer S, Cheng Y, Wong J, Vertatschitsch E, Mate TP. 2005. Accuracy of a wireless localization system for radiotherapy. *Int J Radiat Oncol Biol Phys* 61:933-937.
- Berbeco RI, Neicu T, Rietzel E, Chen GT, Jiang SB. 2005. A technique for respiratory-gated radiotherapy treatment verification with an EPID in cine mode. *Phys Med Biol* 50:3669-3679.
- Brock KK, McShan DL, Ten Haken RK, Hollister SJ, Dawson LA, Balter JM. 2003. Inclusion of organ deformation in dose calculations. *Med Phys* 30:290-295.
- Cai W, Feng C, Schreibmann E, Gambhir S, Wu Y, Xing L, Chen S. 2006. ¹⁸F-Labeled Bombesin Analogs for Targeting GRP Receptor-Expressing Prostate Cancer. *Journal of nuclear medicine*:in press.
- Chang J, Sillanpaa J, Ling CC, Seppi E, Yorke E, Mageras G, Amols H. 2006. Integrating respiratory gating into a megavoltage cone-beam CT system. *Med Phys* 33:2354-2361.
- Chen L, Price RA, Jr., Wang L, Li J, Qin L, McNeeley S, Ma CM, Freedman GM, Pollack A. 2004. MRI-based treatment planning for radiotherapy: dosimetric verification for prostate IMRT. *Int J Radiat Oncol Biol Phys* 60:636-647.
- Coolens C, Evans PM, Seco J, Webb S, Blackall JM, Rietzel E, Chen GT. 2006. The susceptibility of IMRT dose distributions to intrafraction organ motion: an investigation into smoothing filters derived from four dimensional computed tomography data. *Med Phys* 33:2809-2818.
- Coselmon MM, Balter JM, McShan DL, Kessler ML. 2004. Mutual information based CT registration of the lung at exhale and inhale breathing states using thin-plate splines. *Medical physics* 31:2942.
- Court LE, Dong L, Lee AK, Cheung R, Bonnen MD, O'Daniel J, Wang H, Mohan R, Kuban D. 2005. An automatic CT-guided adaptive radiation therapy technique by online modification of multileaf collimator leaf positions for prostate cancer. *Int J Radiat Oncol Biol Phys* 62:154-163.

- de Koste JR, Lagerwaard FJ, de Boer HC, Nijssen-Visser MR, Senan S. 2003. Are multiple CT scans required for planning curative radiotherapy in lung tumors of the lower lobe? *Int J Radiat Oncol Biol Phys* 55:1394-1399.
- Dietrich L, Jetter S, Tucking T, Nill S, Oelfke U. 2006. Linac-integrated 4D cone beam CT: first experimental results. *Phys Med Biol* 51:2939-2952.
- Hawkins MA, Brock KK, Eccles C, Moseley D, Jaffray D, Dawson LA. 2006. Assessment of residual error in liver position using kV cone-beam computed tomography for liver cancer high-precision radiation therapy. *Int J Radiat Oncol Biol Phys* 66:610-619.
- Hector CL, Webb S, Evans PM. 2000. The dosimetric consequences of inter-fractional patient movement on conventional and intensity-modulated breast radiotherapy treatments. *Radiother Oncol* 54:57-64.
- Hugo GD, Agazaryan N, Solberg TD. 2003. The effects of tumor motion on planning and delivery of respiratory-gated IMRT. *Med Phys* 30:1052-1066.
- Jaffray DA, Siewerssen JH, Wong JW, Martinez AA. 2002. Flat-panel cone-beam computed tomography for image-guided radiation therapy. *Int J Radiat Oncol Biol Phys* 53:1337-1349.
- Kamath S, Paquin D, Levy D, Schreibmann E, Xing L. 2006. Incorporating priori knowledge into deformable image registration. In: 2006 Annual Meeting of ASTRO. Philadelphia, PA.
- Langen KM, Meeks SL, Poole DO, Wagner TH, Willoughby TR, Kupelian PA, Ruchala KJ, Haimerl J, Olivera GH. 2005. The use of megavoltage CT (MVCT) images for dose recomputations. *Phys Med Biol* 50:4259-4276.
- Li T, Schreibmann E, Thorndyke B, Tillman G, Boyer A, Koong A, Goodman K, Xing L. 2005. Radiation dose reduction in four-dimensional computed tomography. *Med Phys* 32:3650-3660.
- Li T, Schreibmann E, Yang Y, Xing L. 2006. Motion correction for improved target localization with on-board cone-beam computed tomography. *Phys Med Biol* 51:253-267.
- Li T, Xing L, Munro P, McGuiness C, Yang Y, Cao M, Loo B, and Koong A, Four-dimensional cone-beam computed tomography using an on-board imager, *Medical Physics* 33, 3825-33, 2006.

- Litzenberg DW, Balter JM, Hadley SW, Sandler HM, Willoughby TR, Kupelian PA, Levine L. 2006. Influence of intrafraction motion on margins for prostate radiotherapy. *Int J Radiat Oncol Biol Phys* 65:548-553.
- Paquin D, Levy D, Xing L. 2006. Hybrid multistage landmark and deformable image registration. *IEEE Pattern Recognition and Machine Intelligence*:submitted.
- Paskalev K, Ma CM, Jacob R, Price R, McNeeley S, Wang L, Movsas B, Pollack A. 2004. Daily target localization for prostate patients based on 3D image correlation. *Phys Med Biol* 49:931-939.
- Ruan D, Fessler JA, Balter JM, Sonke JJ. 2006. Exploring breathing pattern irregularity with projection-based method. *Med Phys* 33:2491-2499.
- Schreibmann E, Chen GT, Xing L. 2006. Image interpolation in 4D CT using a BSpline deformable registration model. *Int J Radiat Oncol Biol Phys* 64:1537-1550.
- Sonke JJ, Zijp L, Remeijer P, van Herk M. 2005. Respiratory correlated cone beam CT. *Med Phys* 32:1176-1186.
- Sorcini B, Tilikidis A. 2006. Clinical application of image-guided radiotherapy, IGRT (on the Varian OBI platform). *Cancer Radiother*. 10, 252-257, 2006.
- Thilmann C, Nill S, Tucking T, Hoss A, Hesse B, Dietrich L, Bendl R, Rhein B, Haring P, Thieke C, Oelfke U, Debus J, Huber P. 2006. Correction of patient positioning errors based on in-line cone beam CTs: clinical implementation and first experiences. *Radiat Oncol* 1:16.
- Vedam SS, Keall PJ, Kini VR, Mostafavi H, Shukla HP, Mohan R. 2003. Acquiring a four-dimensional computed tomography dataset using an external respiratory signal. *Physics in Medicine and Biology* 48:45.
- Webb S. 2006. Motion effects in (intensity modulated) radiation therapy: a review. *Phys Med Biol* 51:R403-425.
- Wink NM, McNitt-Gray MF, Solberg TD. 2005. Optimization of multi-slice helical respiration-correlated CT: the effects of table speed and rotation time. *Phys Med Biol* 50:5717-5729.
- Xiong L, Viswanathan A, Steward AJ, Haker S, Tempany CM, Chin L, Cormack RA. 2006. Deformable structure registration of bladder through surface mapping. *Medical Physics* 33:1848-1856.

- Xing, L, Thorndyke B, Schreibmann E, Li T, Yang Y, Kim G., Luxton G, Koong, A,
Overview of image guided radiation therapy (IGRT), *Medical Dosimetry* 31, 91-122,
2006.
- Yin FF, Guan H, Lu W. 2005. A technique for on-board CT reconstruction using both
kilovoltage and megavoltage beam projections for 3D treatment verification. *Med Phys*
32:2819-2826.

Figure Captions

Fig. 1. The CT and CBCT images with and without motion for the Catphan-600: (a) pCT in the absence of phantom motion; (b) CBCT in the absence of phantom motion; (c) pCT with moving phantom; and (d) CBCT with moving phantom.

Fig. 2. (a) The calibration curves (Hounsfield number vs relative electron density) for pCT, half fan and full fan mode CBCT; (b) the variation of calibration curves with time for the full fan CBCT.

Fig. 3. HU profiles of pCT and CBCT images (see figure 2) along the A-A line (panel a) and B-B line (panel b).

Fig. 4. HU profiles for three different sized homogeneous cylindrical water phantoms. The diameters for large, medium and small phantoms are 10.8, 16.0 and 26.6 cm, respectively.

Fig. 5. Dose distributions in a transverse slice calculated based on the four sets of CT data shown in figure 2: (a) pCT; (b) CBCT; (c) pCT with a motion; and (d) CBCT with a motion. In all four situations, a $5 \times 5\text{cm}^2$ single field plan was used to irradiate a spherical hypothetical target with a diameter of 5cm located at the phantom center.

Fig. 6. Comparison of the dose profiles along the two orthogonal lines shown in figure 2 for the Catphan-600 phantom: (a) profile along the A-A line; (a) profile along the B-B line.

Fig. 7. Comparison of the target DVHs calculated based on the four sets of CT data shown in figure 2 for the phantom case.

Fig.8 Dose distributions computed based on CT (top row) and CBCT (middle row) in the CIRS pelvic phantom. The overlay of CT- and CBCT-based doses in the target region on an axial slice is shown in the lower-left panel. The DVHs of the target and sensitive structure is shown in the lower-right panel. In the bottom row, the solid and dotted curves represent the results from CT and CBCT-based calculations, respectively.

Fig. 9. CT, CBCT and mCBCT images for the first prostate case: (a) pCT; (b) daily CBCT; (c) checkerboard overlay of CT and CBCT after the deformation registration; and (d) mCBCT.

Fig. 10. Dose distributions in a transverse slice calculated based on the: (a) pCT; (b) CBCT; (c) mCBCT for the first prostate case. DVHs of the prostate, PTV, rectum and bladder obtained based on the pCT, CBCT and mCBCT images for the case is plotted in (d).

Fig. 11. Comparison of DVHs of the prostate, PTV, rectum and bladder obtained based on the pCT, CBCT and mCBCT images for the second prostate case.

Fig. 12. Comparison of DVHs of the prostate, PTV, rectum and bladder obtained based on the pCT, CBCT and mCBCT images for the third prostate case.

Fig. 13. CT, CBCT and modified CT images for the lung case: (a) planning CT; (b) CBCT; (c) checkerboard image after the deformation registration; and (d) modified CBCT.

Fig. 14. Dose distribution in a transverse slice calculated based on: (a) planning CT; (b) CBCT; and (c) modified CBCT for the lung case. A comparison of the DVHs of GTV, PTV, right lung and spinal cord obtained based on planning CT, CBCT and the modified CBCT images is shown in (d).

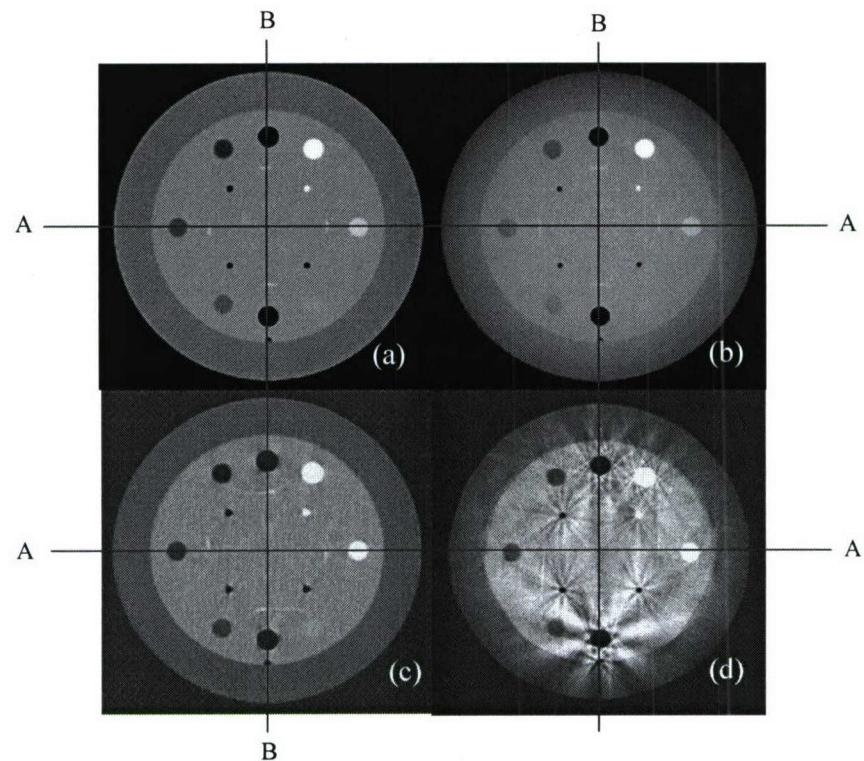
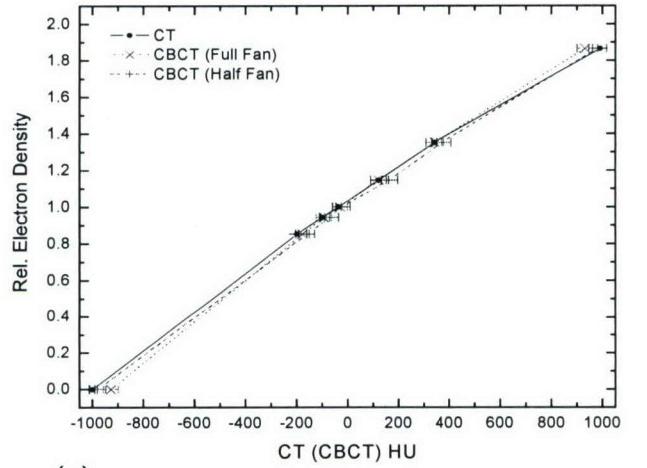
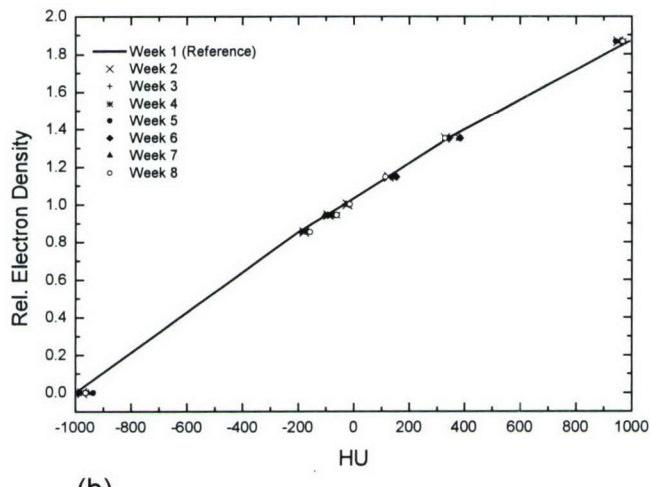


Figure 1

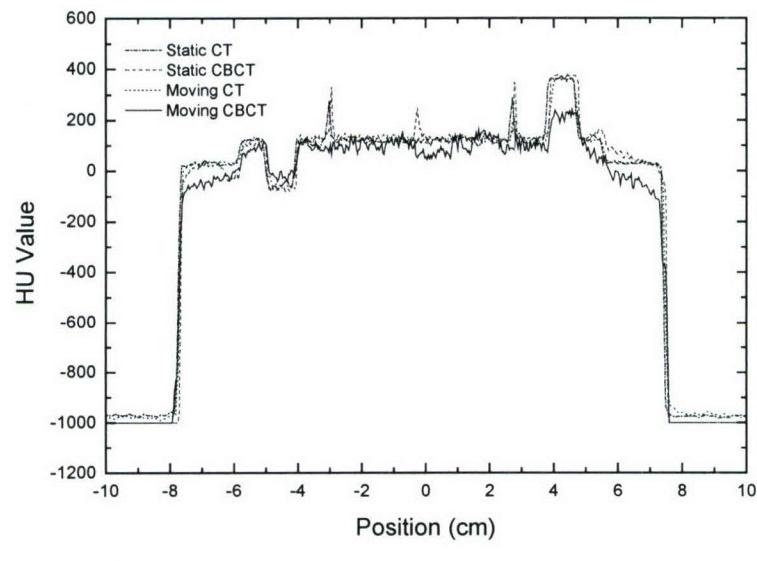


(a)

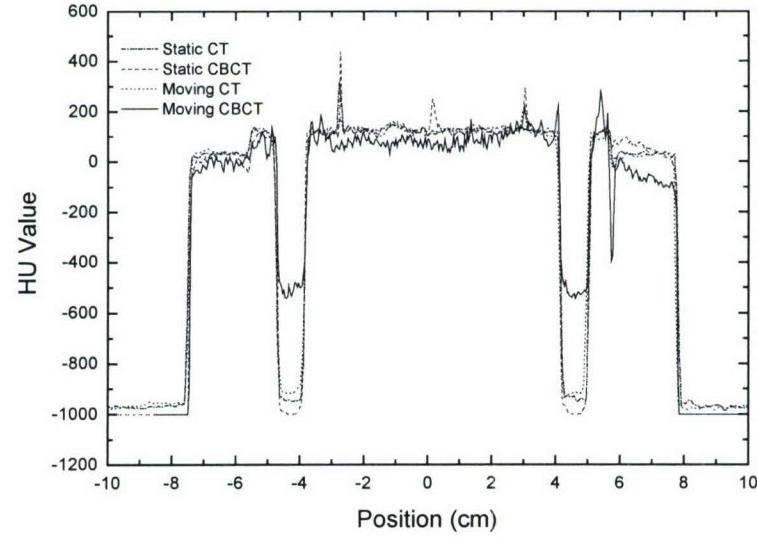


(b)

Figure 2



(a)



(b)

Figure 3

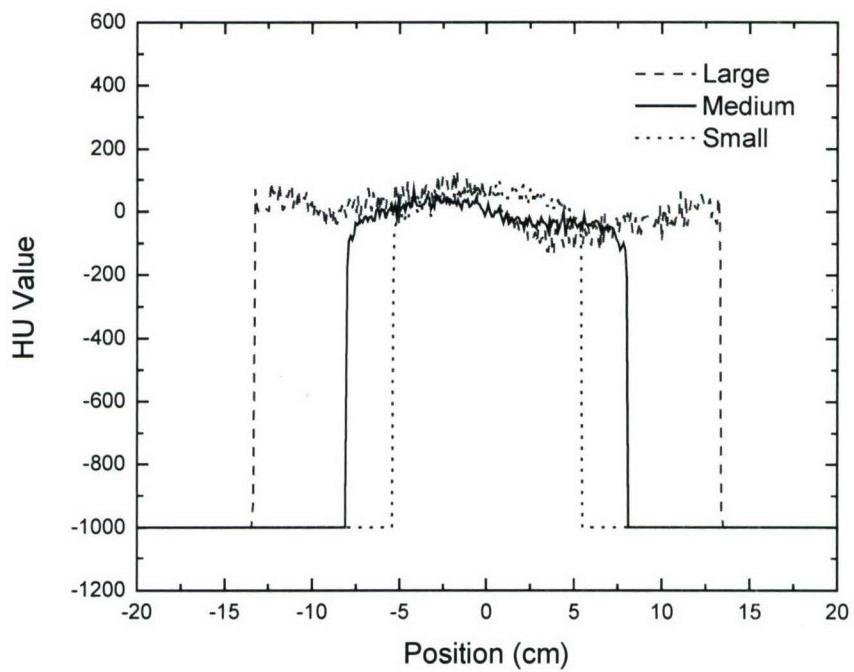


Figure 4

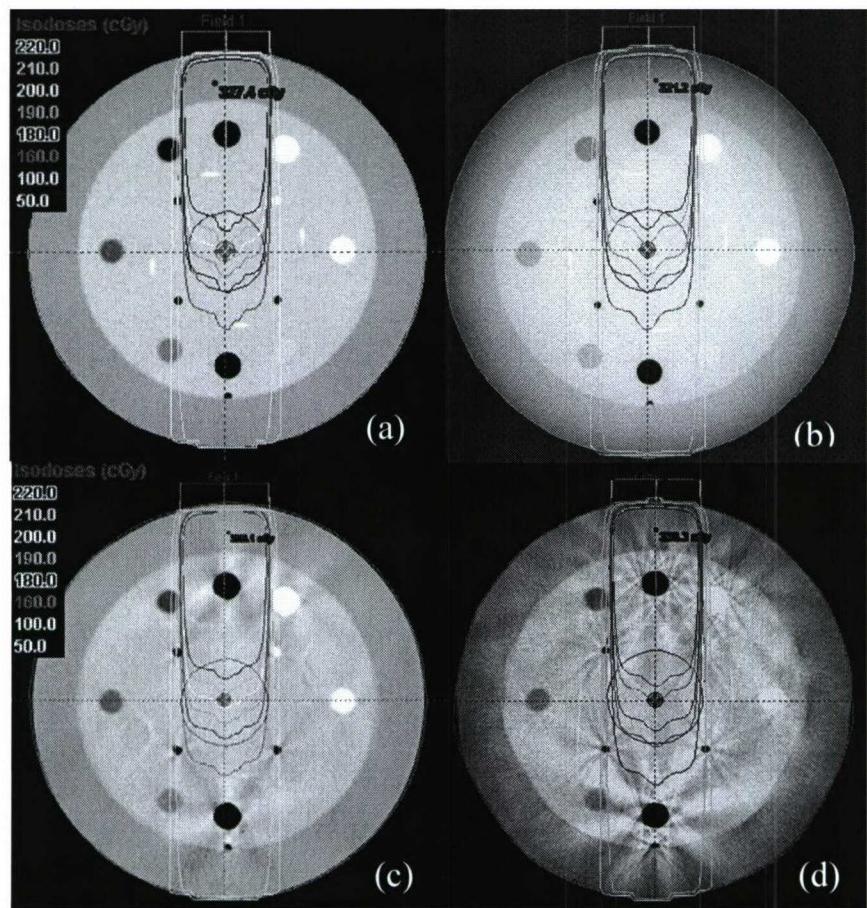
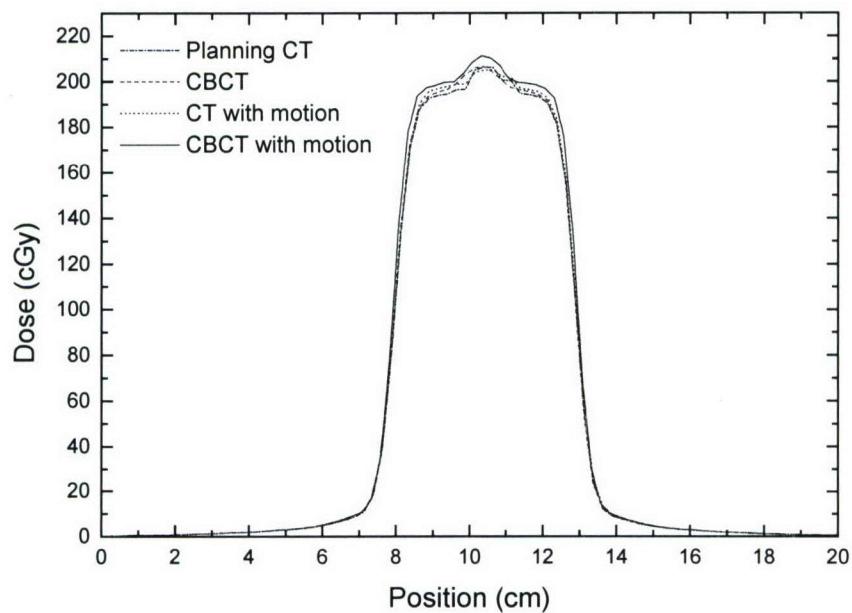
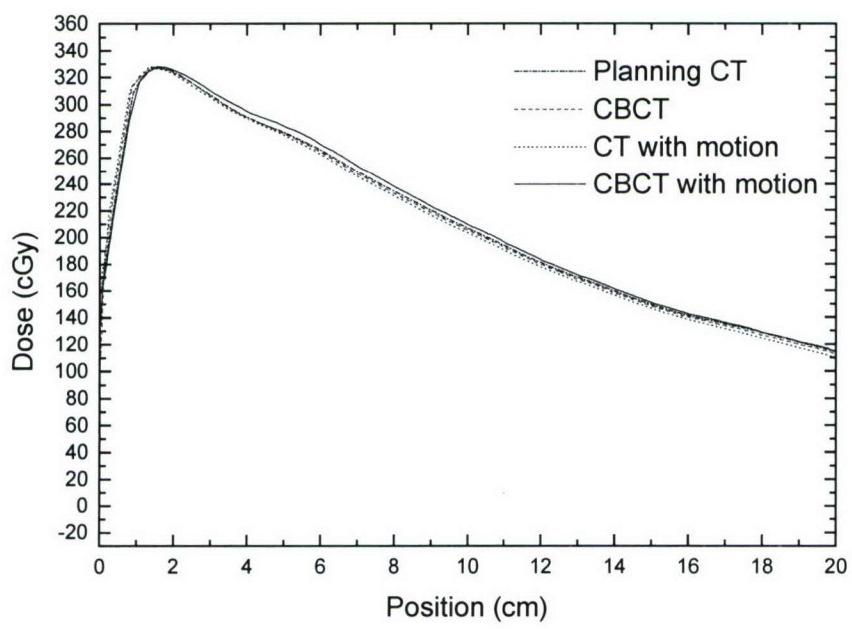


Figure 5



(a)



(b)

Figure 6

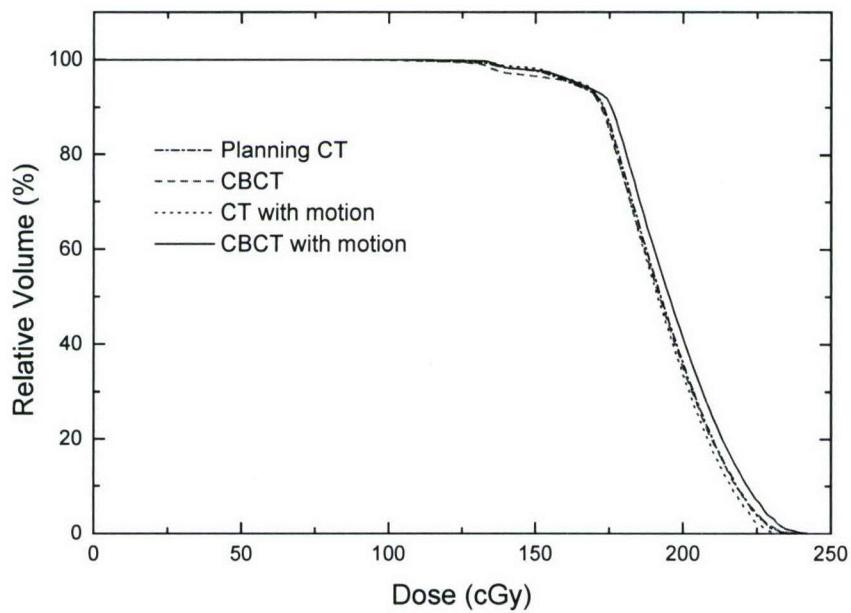


Figure 7

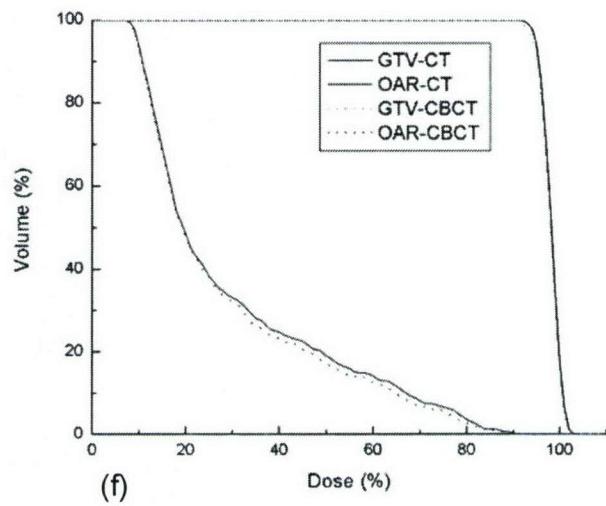
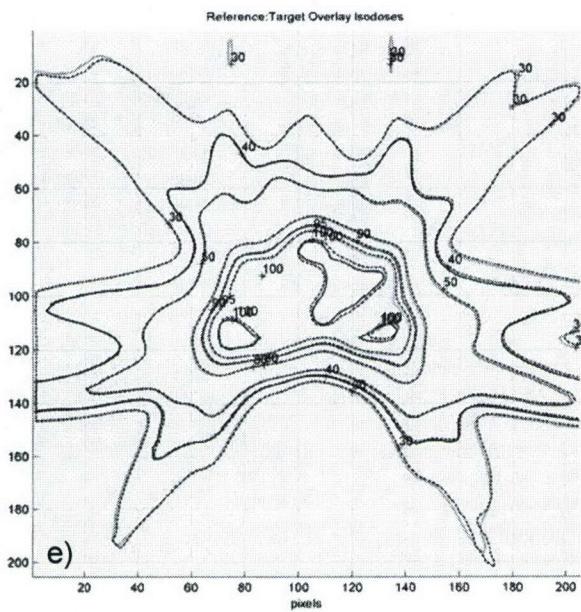
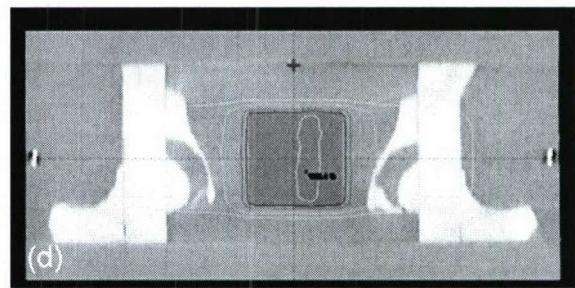
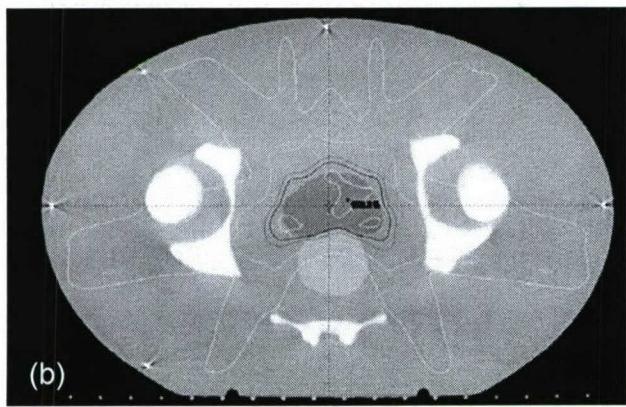
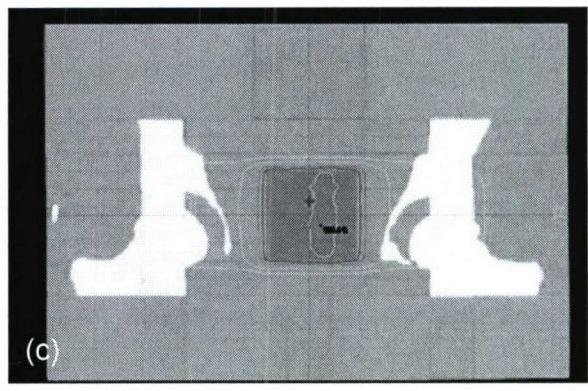
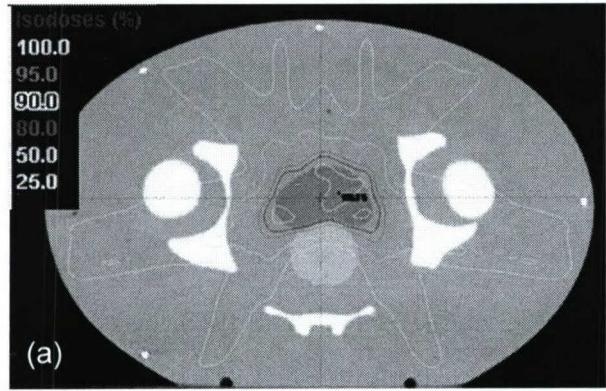


Figure 8

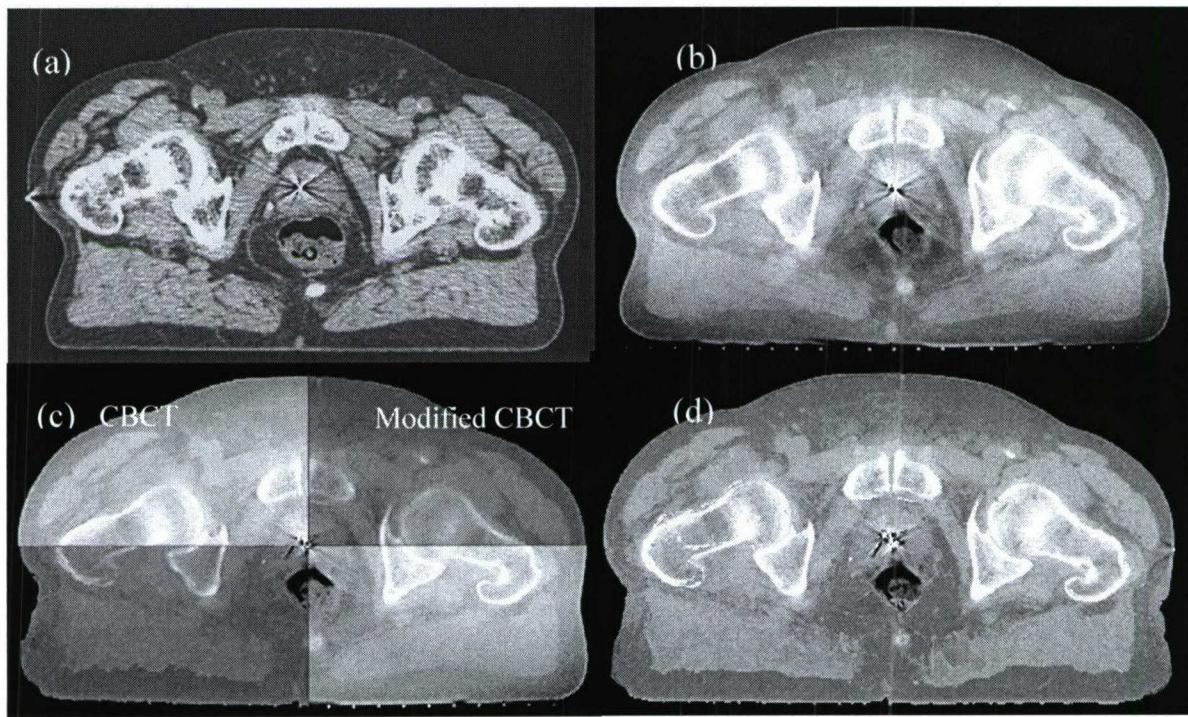


Figure 9

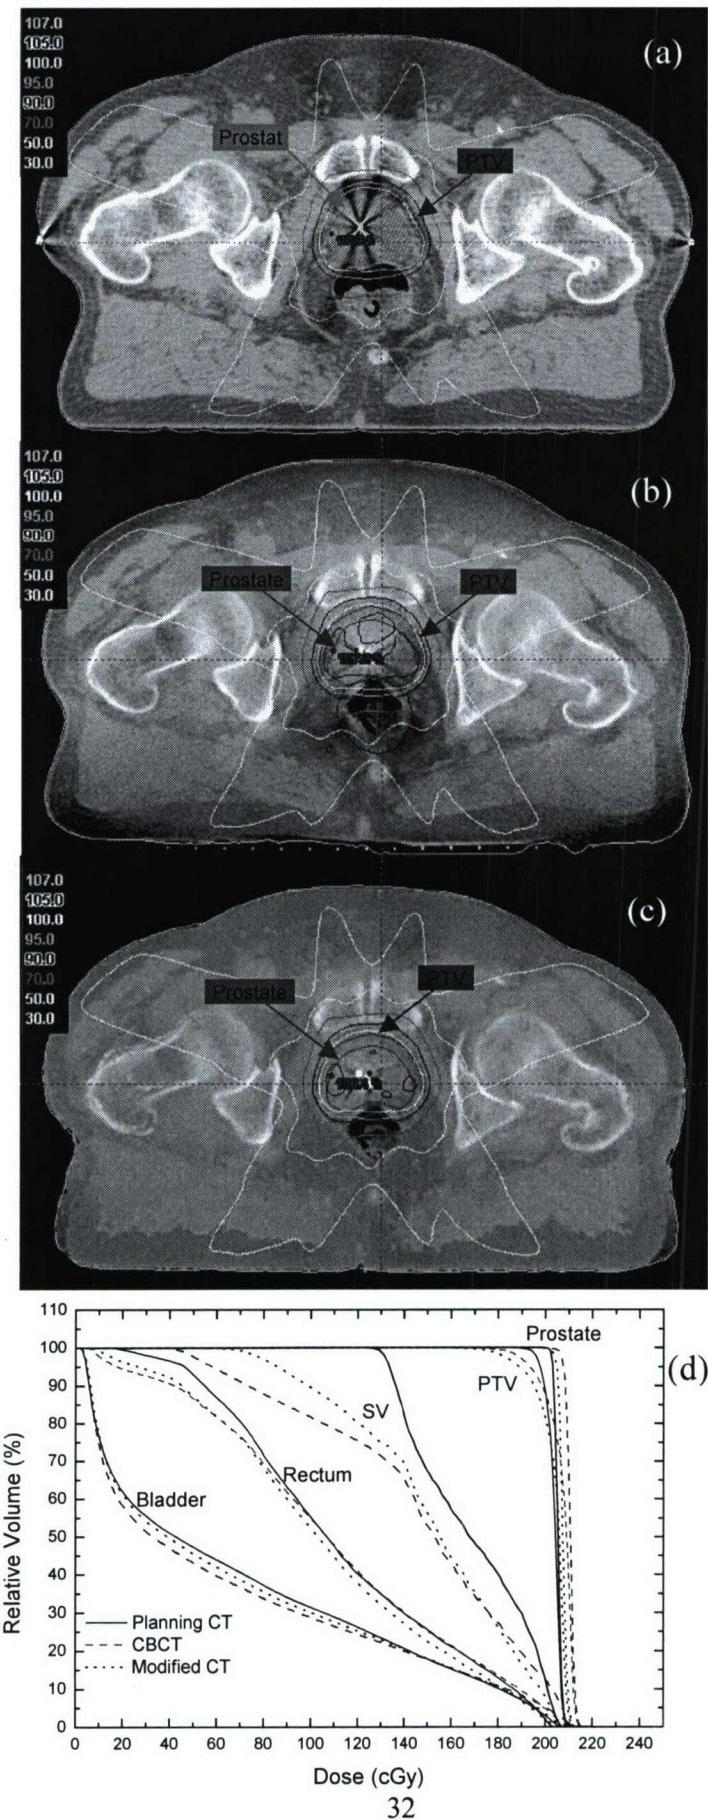


Figure 10

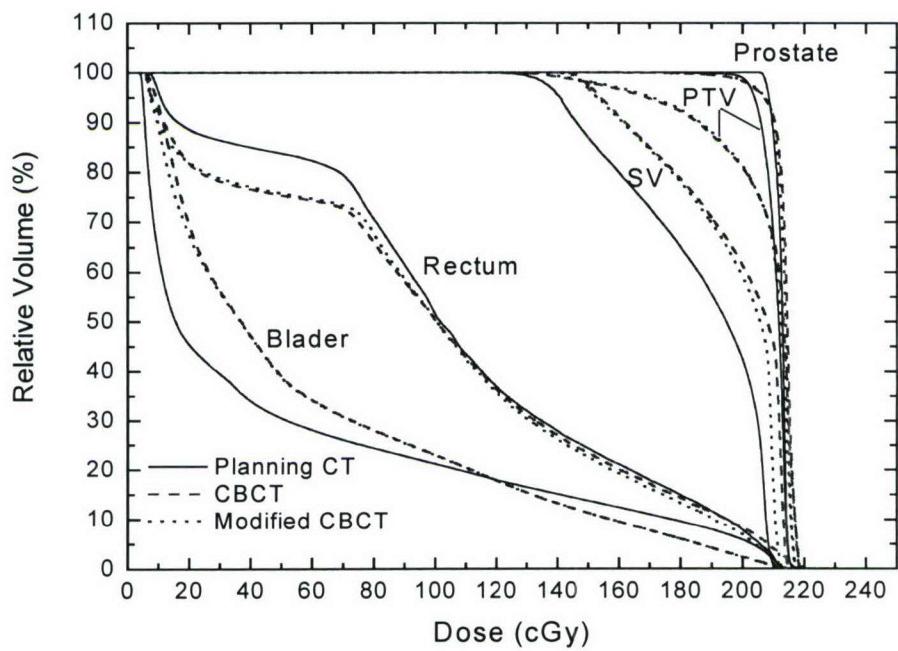


Figure 11

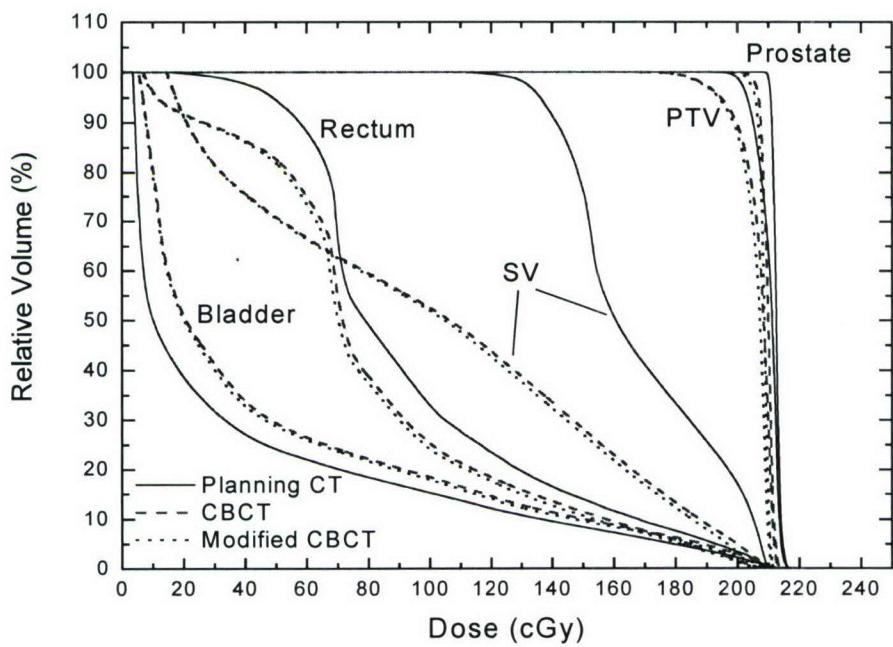


Figure 12

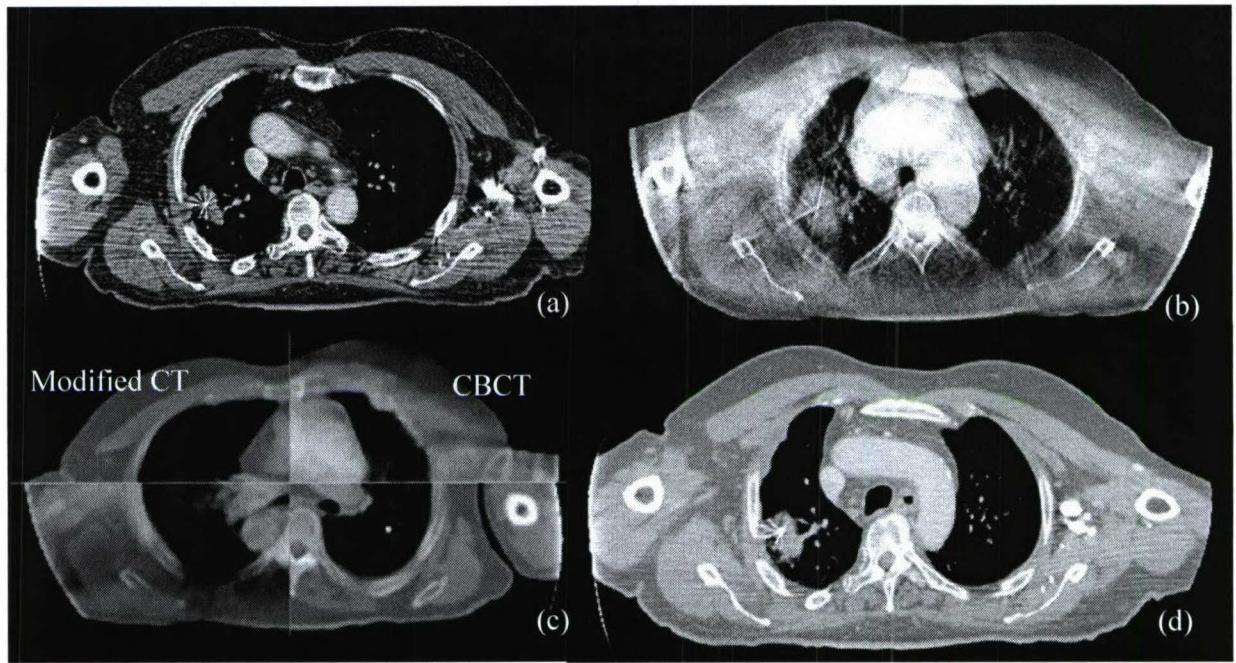
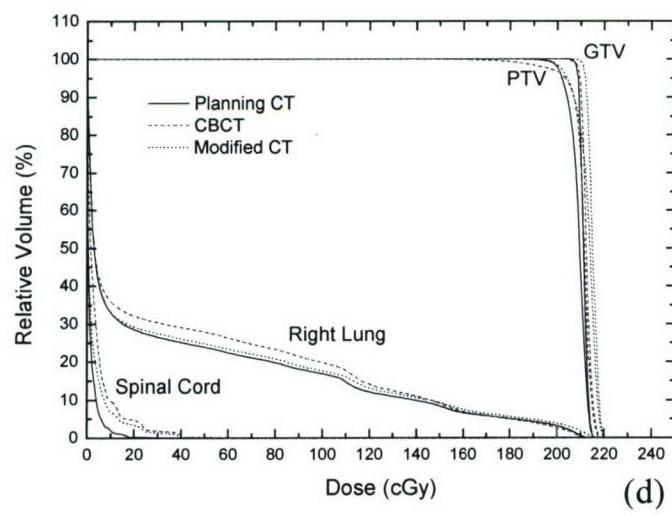
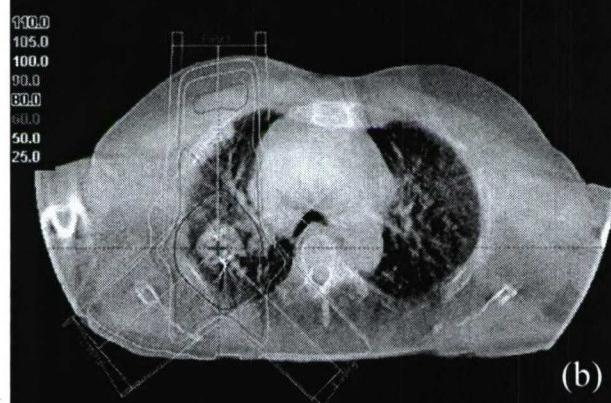
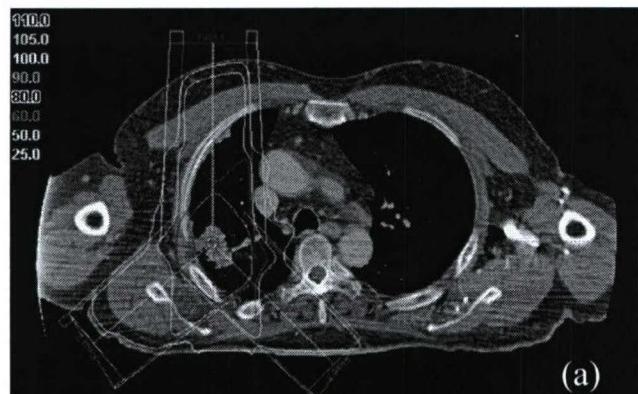


Figure 13



Automated Contour Mapping for 4D Radiation Therapy

5

Ming Chao, Ph.D., Eduard Schreibmann, Ph.D., Tianfang Li, Ph.D., Nicole Wink, Ph.D.,
Lei Xing^{a)}, Ph.D.

10

Department of Radiation Oncology, Stanford University School of Medicine, Stanford,
California 94305-5847

15

a) Author to whom correspondence should be addressed.

Stanford University School of Medicine

Department of Radiation Oncology

20 875 Blake Wilbur Drive

Stanford, CA 94305-5847

E-mail: lei@reyes.stanford.edu

Phone: (650)498-7896

Fax: (650)498-4015

25

Submitted to: *Medical Physics*

ABSTRACT

The purpose of this work is to develop a novel strategy to automatically map organ contours from one phase of respiration to all other phases on a four-dimensional computed tomography (4D CT). A region of interest (ROI) was manually delineated by a physician on one phase specific image set of a 4D CT. A number of cubic control volumes of size of ~1cm were automatically placed along the contours. The control volumes were then collectively mapped to the next 4D CT phase set using a rigid transformation. To accommodate the organ deformation, a model-based adaptation of the control volume positions was followed after the rigid mapping procedure. This further positional adjustment of control volume positions was performed by minimizing an energy function which balances the tendency for the control volumes to move to their correspondences with the desire to maintain similar image features and shape integrity of the contour. The mapped ROI surface was then constructed based on the central positions of the control volumes using a triangulated surface construction technique. The proposed technique was assessed using 4D CT images of a digital phantom and three lung patients. Our digital phantom study data indicated that a spatial accuracy better than 2.7 mm is achievable using the proposed technique. The patient study showed a similar level of success. In addition, the computational speed of our algorithm was significantly improved as compared with a conventional deformable registration-based contour mapping technique. The robustness and accuracy of this approach make it a valuable tool for the efficient use of the available spatial-tempo information for 4D simulation and treatment.

55 **Key words:** CT, image registration, contour mapping, IGRT

I. INTRODUCTION

A longstanding question in radiation therapy is how to accurately and efficiently segment
60 a region of interest (ROI) such as a tumor target volume or a sensitive structure¹⁻⁶. In
spite of intense research efforts in the past few decades, ROI segmentation remains a time
consuming task in treatment planning. In most cases, the segmentation is performed
manually in a slice-by-slice fashion, creating a strong need for automated segmentation
tools in the clinics. The introduction of four-dimensional computed tomography (4D CT)
65 in radiation oncology practice further amplifies this need as the number of images to be
segmented is increased dramatically⁷⁻¹⁵. Generally, a 4D CT scan consists of 5-10 sets of
three-dimensional (3D) CT images, each representing the patient anatomy at a specific
phase of respiration. For 4D radiation therapy applications, it is impractical to follow the
3D approach of manual segmentation due to the immense workload associated with this
70 process.

A natural way to deal with the 4D segmentation problem is to start with a known
set of contours for a selected phase and map these contours onto all other phases. ROIs
for the selected phase are first manually contoured, similar to that done in treatment
planning based on 3D CT image data sets. The mapping procedure can be accomplished
75 with the aid of a computer algorithm that registers an arbitrary point on the selected phase
to the corresponding points on all other phases. While conceptually simple, the
implementation of this idea is not straightforward. An intelligent algorithm capable of
providing accurate point-to-point correspondence between the phased images, or at least
between points within the ROIs, is the key to the success of this approach. Various
80 studies have investigated algorithms to automatically map contours using deformable
image registration and surface mapping techniques. One method is based on a deformable
registration model¹⁶⁻¹⁸. This method has limited accuracy, especially in the regions
proximate to the interfaces of different organs, and is brute-force in nature, which entails
a large amount of computations. In reality, contour mapping is a regional problem and a
85 global association of the phase-based images is neither necessary nor efficient. Surface
mapping achieves the stated goal of contour transformation by iteratively deforming the
ROI contour-extended surface until the optimal match with the reference is reached^{2, 19-22}.
Numerous surface mapping techniques, such as spatial partitioning, principal component

analysis, conformal mapping, rigid affine transformation, deformable contours and
90 warping based on the thin-plate spline (TPS), have been developed over the years and the endpoint of all of these techniques is a mapping between topological components of the input surfaces that allow for transfer of annotations. This type of computation is inherently more efficient in comparison with the deformable model-based approaches, but it suffers from the fact that the resultant mapping heavily depends on the model used
95 and the fact that the model parameters in the calculations are not physically transparent.

In this work we developed a hybrid approach to contour mapping that combines the useful features of the two different types of techniques. Our working hypothesis is that information contained in the boundary region is often sufficient to guide the contour mapping process without relying on the use of an *ad hoc* surface deforming model. In the
100 proposed technique, the neighborhood information of the contour surface is captured by a series of small cubic (or other shaped) control volumes placed around the surface²³. The collection of the centers of the control volumes is a representation of the contour surface. The ROI contour mapping proceeds iteratively under the guidance of the information contained in the control volumes. The proposed approach is conceptually simple and
105 straightforward to implement. Application of the method to a digital phantom and three clinical lung cases suggests that the technique is quite robust and could make a significant impact in 4D radiation therapy.

110 II. MATERIALS AND METHOD

A. Software platform

The Insight Toolkit (ITK)²⁴ and the Visualization Toolkit (VTK)²⁵, which are open source cross-platform C++ software toolkits and are freely available for research
115 purposes (<http://www.itk.org> for ITK and <http://public.kitware.com/VTK/> for VTK), were used in this study. A variety of methods have been programmed into the ITK platform for image registration and segmentation. ITK was used for automatic mapping while VTK was mainly used for 3D visualization and contour construction.

120 **B. Image acquisition**

The 4D CT image data sets for three lung cancer patients were acquired with a multislice helical CT scanner (Discovery ST, GE Medical System, Milwaukee, WI). The collected data was sorted into ten phase bins¹⁰. The 4D CT image sets for all patient studies were reconstructed with a 2.5 mm slice thickness. The size and pixel resolution of
125 each CT slice was 512×512 and 0.98×0.98 mm², respectively. The 4D CT images were transferred through DICOM to a personal computer (PC) with a Pentium IV (2.66 GHz) processor for image processing. One of the phases, which is referred to as the template phase, was selected for manual segmentation of the ROIs. We refer to the other phases in the 4D CT image set as the target phases with corresponding target contours.

130

C. Placement of control volumes along the ROI contour surface

The task of mapping a contour is to find its corresponding location on the target phase for an arbitrary point on the contour drawn on the template image. In general, the image feature surrounding a contour point can be used as a signature of the point to aid
135 the search for its corresponding position on the target phase. In this study, the image feature at each point was captured by introducing a cubic control volume (~1cm in size) centered at the point. After the manual ROI delineation was performed on the template phase, a number of cubic control volumes were automatically placed along the ROI contours. This is depicted in figure 1, where the green curve represents the contour and
140 the pink squares represent the control volumes. The center of each control volume was set at a contour point, typically 0.5~1cm from the next point, as illustrated in figure 1. The collection of these points represents the ROI contour surface.

The contour mapping was carried out in three steps: (i) mapping the introduced control volumes collectively using a rigid image registration algorithm; (ii) iteratively
145 fine-tuning the 3D positions of the control volumes to determine the deformed ROI contour surface; and (iii) reconstructing the new contours by cutting the deformed surface slice-by-slice along the transversal, sagittal or coronal direction. This procedure is shown in figure 2 and described in detail below.

150 **D. Collective mapping of control volumes**

After a series of control volumes were placed along the segmented contours on the template phase, we mapped them onto the target phase collectively (i.e., all the control volumes were treated as an entity) using a rigid image registration algorithm²⁴. A feature of the rigid collective mapping is that the relative distances and orientations of the control volumes remain the same during the course of mapping, resulting in an approximate ROI contour surface in the target phase and providing a good start for further adjustment of the contour shape to accommodate the organ deformations. We note that using a rigid mapping is not a necessary step and other methods capable of providing a reasonable initial estimate of the ROI surface may also be used.

160 The collective mapping of the control volumes was performed under the guidance of a normal cross correlation (NCC) metric defined by

$$f = -\frac{\sum_{\alpha} \sum_{\beta} \sum_{i=1} I_{\alpha}(\mathbf{x}_i) I_{\beta}(\mathbf{x}'_i)}{\sqrt{\sum_{\alpha} \sum_{i=1} I_{\alpha}^2(\mathbf{x}_i) \sum_{\beta} \sum_{j=1} I_{\beta}^2(\mathbf{x}'_j)}}, \quad (1)$$

where $I_{\alpha}(\mathbf{x}_i)$ is the intensity at a point \mathbf{x}_i in a control volume indexed by α in the target phase, and $I_{\beta}(\mathbf{x}'_i)$ represents the intensity at a point \mathbf{x}'_i in a control volume indexed by β in the template phase. In Eq. (1), \mathbf{x} and \mathbf{x}' are related by a rigid transformation \mathbf{T} , where $\mathbf{T}\mathbf{x}' = \mathbf{x}$. We calculated the transformation matrix \mathbf{T} , which maps the ensemble of control volumes from the template phase to the target phase. The limited memory Broyden-Fletcher-Goldfarb-Shannon algorithm (L-BFGS) algorithm²⁶⁻³⁰ was used to optimize the metric in Eq. (1) with respect to the transform parameters. The details of this algorithm have been described elsewhere^{18, 31} and will not be repeated here. During the course of the control volume mapping, an iterative calculation based on the L-BFGS algorithm was carried out until a preset maximum number of iterations were met or until the NCC function stopped improving.

175 **E. Fine tuning of the mapped control volumes**

In the absence of deformation, the ROI contour of the target phase is obtained by connecting the centers of the rigidly mapped control volumes. In a more general case

where deformation does exist, the contour from the rigid mapping serves as an initial estimate of the ROI contour in the target phase. Further positional adjustment of the control volumes is needed to accommodate the deformation of the ROI. The final positions of the control volumes, which define the ROI surface, are determined by balancing the tendency for the control volumes to move to their corresponding locations with the desire to maintain similar image features and shape integrity of the contour.

Mathematically, the above adaptation process is modeled by two “energy terms”.
185 The first term is referred to as “self-energy” and the second term describes the “interaction” among the control volumes. Self-energy tends to drive a control volume toward a position where the neighborhood environment resembles itself most. For each control volume in the template image, this process was driven by the NCC between the volume and its corresponding locations in the target image. The interaction term among
190 the control volumes intends to maintain the integrity of the control volume cluster as a whole and prevents any unrealistic control volume configuration from happening. The interaction energy term is described by

$$E_{\text{interaction}} = M_{\text{template}} - M_{\text{target}}, \quad (2)$$

where M_{template} is defined as the correlation between a control volume in the template phase and the neighboring control volumes in the target phase, M_{target} represents the correlation between the control volume in the target phase and the neighboring volumes in the template phase. These two terms take the neighborhood environment of the control volume being adjusted into account and apply a constraint on the possible form of the control volume configuration. In a sense, the interaction energy in Eq. (2) exerts a restoring force when the position of a control volume is varied with respect to its neighbors. The interaction force is important in preventing the control volumes from moving to unrealistic positions simply driven by the self-energy and in retaining the shape integrity of the ROI surface. For simplicity, only the adjacent control volumes were considered when computing the interaction energies. The final position of each control
200 volume was determined by minimizing the sum of the self-energy and interaction energy terms.
205

A simple searching algorithm was implemented to find the minimum of Eq. (2) and thus the final configuration of the control volumes. For each rigidly mapped control

volume on the target phase, we defined a small region of ~3 cm around the central point
210 in the volume. Eq. (2) was then minimized to determine the optimal control volume location within this region. This process was repeated for each separate control volume. Due to the fact that the interaction energy term demands the information from the neighboring volumes, adjusting the location of each control volume changes the position of the previously fine-tuned volume. Therefore, several cycles of this process were
215 needed to obtain the truly optimal positions of all the control volumes. We found (see digital phantom study in results section) that 2 - 3 complete adjustment cycles were adequate to find the optimal control volume configuration.

F. Reconstruction of target contours

Upon completion of the control volume mapping and adaptation, the centers of each control volume were identified. An ROI surface was constructed with these central points using a triangulated surface construction technique. The intersection of the surface with each CT slice was superimposed on top of the image in order to visualize the contour in a conventional fashion. These contours can be exported as ASCII or DICOM-
225 RT format for treatment planning.

G. Evaluation of the algorithm and case study

Evaluation of a contour mapping algorithm is a difficult task because of the general lack of a gold standard for comparison. The proposed control volume based 4D contouring technique was first evaluated with a digital phantom experiment. In this study, a thoracic CT image was deformed intentionally using a known deformation matrix (the deformation matrix was introduced by drifting the positions of each control volume along the contour with known sizes. This way the CT image was deformed because of each voxel's positional change). The "ground truth" lung surface in the target phase was attainable by transforming the original contours with the deformation matrix. The manually delineated lung contours in the original image were also mapped using the proposed novel control volume based approach. A comparison of the mapped lung surface with the ground truth allowed us to quantitatively assess the success of the proposed approach.
230
235

240 The control volume based contour mapping technique was also applied to three
4D CT patient scans. A physician manually delineated the lungs and GTV for each scan
on a selected phase (for example, the exhale phase). These contours were then mapped
onto all other phases using the proposed technique. Visual inspection was used to
evaluate the three patient studies.

245

III. RESULTS

A. Digital phantom study

250 The thoracic CT images before and after the intentionally introduced deformation
are shown in Fig. 3. The manually delineated lung contour is shown in green in figure 3a,
while the “ground truth” contour obtained by transforming the manual contour with the
known transformation matrix is plotted in figure 3b in yellow. The contour from the
proposed approach is shown in red in the same figure. To be comprehensive, the rigidly
mapped contour, which served as the start of the model-based refinement, is also plotted
in figure 3b in blue. Visual inspection of the yellow and red contours in figure 3b
255 indicated that the gold standard and the mapped contours are similar, despite the fact that
the intentionally introduced deformation field is quite large (the displacement of some
of the voxels is as large as 2.0 cm). The mean and maximum separations between the two
sets of contours were found to be 1.5 mm, and 2.7 mm, respectively.

260 To better understand the mapping process, we examined the behavior of each
energy term (metrics) during the course of the contour mapping. In figure 4a we plotted
the values of self-energy, interactive energy and total energy of one of the control
volumes in the process of positional adjustment in the surrounding area of the control
volume in the deformed image. For this particular control volume, the minimum of Eq. (2)
was reached after searching the first 30 points. For control volumes located in regions
265 where the deformation is large, more searching points may be required to reach the
minimum. It is also interesting to show how the average metric value of all the control
volumes evolved. In figure 4b the metric value is depicted as a function of the calculation
process. The first point, corresponding to step 0, is the metric before mapping. Step 1
refers to the system after rigid mapping. The metric at this point is decreased but does not
reach the minimum. The third point shows the metric after each of the control volumes is
270

fine tuned sequentially. To obtain the optimal contours two more cycles of fine tuning were performed; however, the improvement with further iterations was not significant.

B. Patient studies

Figure 5 shows the manually delineated contours on the 4D CT scans of three lung cancer patients. Contours were drawn on the exhale phases (phase 1) for patients 1 and 2 (figures 5a and b, respectively). In these figures bodies are contoured in red and the lungs are in green. The contours for patient 3 were delineated on the inhale phase (phase 5) as shown in figure 5c. The body contour is in red, lung in yellow, PTV in purple and GTV in green. For patients 1 and 2 the mapping of the lung contours was studied. For the third patient, GTV contour mapping was investigated.

The contours before and after the control volume-based mapping for the first and second patients are presented in figures 6 and 7, respectively. To better visually evaluate the results, three different axial slices are shown for three target phases (phases 4, 7, and 10 for the first patient and phases 2, 5 and 8 for the second patient). In the target phases shown in the middle and left columns, the template contours before and after rigid mapping are shown in green and blue, respectively. The final contours after the model-based adaptation are depicted in red.

When the lung deformation is small (for example, in phase 10 of patient 1 or phases 2 and 8 of patient 2), the rigidly mapped contours closely resemble the target contours. Fine-tuning merely provides limited adjustment of the contours. For phases with large deformations (e.g., phase 4 of patient 1 or phase 5 of patient 2), the model-based adjustment after the rigid contour mapping plays an important role in obtaining optimal contour arrangement. For the lung contours, our analyses of the patient data indicate that the accuracy of our contour mapping technique is within 0.5mm in the superior part of the lung, where respiration induced deformation is small (see (c), (f) and (i) of figures 5 and 6). For other slices, the average error on the target contours is estimated to be less than 2.5 mm.

The mapped GTV contours for phases 2, 4, 7 and 9 in the study of case 3 are shown in figure 8. For phases 4 and 7 the deformation was relatively small and our contour mapping algorithm performed reliably. For phases 2 and 9 which corresponded

to the exhale phases, significant deformation in the ROIs was observed. Despite this less than ideal case, our algorithm still worked well in these phases.

305

IV. DISCUSSION

As in 3D radiation therapy, delineation of ROIs in 4D CT images is a necessary step for treatment planning³². 4D segmentation is required for constructing a 4D patient model and for computing the accumulated dose of moving or deformed organs. Conventionally, in 3D radiation therapy contours are manually drawn in a slice-by-slice fashion; however, an extension of this approach to 4D radiation therapy becomes impractical because of the dramatically increased number of images. One way to tackle this problem is to draw all contours on one of the phases of respiration and map or propagate these contours onto the remaining phases. Various deformable models, which establish a voxel-to-voxel correspondence between different phases, have been used for ROI contour mapping. In this work, we do not use a global registration approach, but rather a regional approach based on the mapping and adaptation of the control volumes. Our approach takes advantage of the imaging features surrounding the ROI and uses them as a guidance in searching for the optimal mapped contours while considering the shape integrity of the ROI surface.

The mapping of a point in one image to another is easily achievable if a unique identifier or signature can be tagged to the point. In our approach, the image feature contained in a control volume is employed as a signature of the point to facilitate the process of finding its corresponding location in the target phases. The concept of control volume was first introduced by Schreibmann and Xing²³ and its advantages for both intra- and inter-modality image registration have been demonstrated. This study represents a novel application of the concept to 4D image segmentation. The central idea is that control volumes placed along the ROI surface carry the neighborhood information of the ROI surface and can be employed to establish a reliable association between the ROIs in two phase specific image sets. After a rigid mapping of the contours from the template phase to the target phase, a model-based adaptation is performed. This

adaptation takes into account the deformation of an object to ensure the overall integrity of the resultant ROI surface. The calculation is local in nature, which improves both
335 computational efficiency and convergence behavior.

A common problem in image segmentation and contour mapping studies is the lack of quantitative validation. In the study of Lu *et al*³³, for example, the accuracy of a deformable model-based contour mapping technique was evaluated purely based on visual inspection. Although it is less quantitative, visual inspection is a convenient way
340 for rapid assessment of a segmentation calculation, especially in a case where the “ground truth” contours do not exist. The same approach was employed in many other previous investigations^{4, 6, 16, 18}. In our study, we introduced a set of digital phantom experiments for evaluating the success of the proposed technique. The “ground truth” ROI surface in this case is known by applying a pre-specified deformation matrix to the
345 original image and companying contours. The digital experiments, along with the three patient studies, suggest that the proposed approach can perform very well even in the presence of significant deformations.

350 **V. CONCLUSION**

The development of 4D radiation therapy involves the use of a large number of images acquired at different times and/or with different modalities. Clinical implementation of the new IGRT paradigm is, to a large extent, bottlenecked by the
355 inability to accurately and efficiently register images and segment the regions of interest (ROIs). In this work, a mathematical framework for control volume-based contour mapping has been proposed for 4D radiation therapy. We demonstrated that the information contained in the boundary region is sufficient to guide the contour mapping process without registering the whole image or relying on the use of an *ad hoc* surface
360 deforming model. The results showed that the control volume-based contour mapping algorithm is capable of robustly and accurately mapping contours from one phase of a 4D CT to the remaining phases. Our technique decreases the workload involved in 4D CT

ROI segmentation and provides a valuable tool for the efficient use of available spatial-tempo information for 4D simulation and treatment.

365

ACKNOWLEDGEMENTS

This work was supported in part by grants from the National Cancer Institute (1R01 CA98523 and CA104205) and Department of Defense (PC040282).

370

1. Kass MR, Witken A, Terzopoulos D. Snakes: active contour models. *Int J Comput Vis* 1988;4:321-331.
2. Coote T, Hill A, Taylor C, Haslam J. The use of active shape models for locating structures in medical images. *J Image Vis Computation* 1994;12:355-366.
3. Xu C, Prince JL. Snakes, shapes, and gradient vector flow. *IEEE Trans. Image Process.* 1998;7:359-369.
4. Liu F, Zhao B, Kijewski PK, Wang L, Schwartz LH. Liver segmentation for CT images using GVF snake. *Med Phys* 2005;32:3699-3706.
5. Weese J, Kaus MR, Lorenz C, al. e. Shape constrained deformable models for 3D medical image segmentation. 2001;Vol. 2082:380-387.
6. Pekar V, McNutt TR, Kaus MR. Automated model-based organ delineation for radiotherapy planning in prostatic region. *Int J Radiat Oncol Biol Phys* 2005;60:973-980.
7. Ritchie CJ, Hsieh J, Gard MF, Godwin JD, Kim Y, Crawford CR. Predictive respiratory gating: a new method to reduce motion artifacts on CT scans. *Radiology* 1994;190:847-852.
8. Vedam SS, Keall PJ, Kini VR, Mostafavi H, Shukla HP, Mohan R. Acquiring a four-dimensional computed tomography dataset using an external respiratory signal. *Phys Med Biol* 2003;48:45-62.

9. Xing L, Thorndyke B, Schreibmann E, Yang Y, Li TF, Kim GY, Luxton G, Koong A. Overview of image-guided radiation therapy. *Med Dosim* 2006;31:91-112.
- 395 10. Pan T, Lee TY, Rietzel E, Chen GT. 4D-CT imaging of a volume influenced by respiratory motion on multi-slice CT. *Med Phys* 2004;31:333-340.
11. Dietrich L, Jetter S, Tucking T, Nill S, Oelfke U. Linac-integrated 4D cone beam CT: first experimental results. *Phys Med Biol* 2006;51:2939-2952.
- 400 12. Li T, Schreibmann E, Thorndyke B, Tillman G, Boyer A, Koong A, Goodman K, Xing L. Radiation dose reduction in four-dimensional computed tomography. *Med Phys* 2005;32:3650-3660.
13. Lu W, Mackie TR. Tomographic motion detection and correction directly in sonogram space. *Phys Med Biol* 2002;47:1267-1284.
- 405 14. Li T, Xing L, Munro P, McGuinness C, Chao M, Yang Y, Loo B, Koong A. Four-dimensional cone-beam computed tomography using an on-board imager. *Med Phys* 2006;33:3825-3833.
15. Sonke JJ, Zijp L, Remeijer P, van Herk M. Respiratory correlated cone beam CT. *Med Phys* 2005;32:1176-1186.
- 410 16. Guerrero TM, Zhang G, Huang TC, Lin KP, Giraud P, De Rycke Y, Dubray B, Helfre S, Voican D, Guo L, Rosenwald JC, Keraudy K, Housset M, Touboul E, Cosset JM. Intrathoracic tumour motion estimation from CT imaging using the 3D optical flow method. *Phys Med Biol* 2004;49:4147-4161.
17. Ragan D, Starkschall G, McNutt T, Kaus M, Guerrero T, Stevens CW. Semiautomated four-dimensional computed tomography segmentation using deformable models. *Med Phys* 2005;32:2254-2261.
- 415 18. Schreibmann E, Chen GT, Xing L. Image interpolation in 4D CT using a BSpline deformable registration model. *Int J Radiat Oncol Biol Phys* 2006;64:1537-1550.
19. McInerney T, Terzopoulos D. Deformable models in medical image analysis. *Med Image Anal* 1996;1:91-108.
- 420 20. Chakraborty A, Staib LH, Duncan JS. An integrated approach for surface finding in medical images. IEEE Workshop Mathematical Methods in Biomedical Image Analysis; 1996. pp. 253-262.

21. Montagnat J, Delingette H, Ayache N. A review of deformable surfaces: topology, geometry and deformation. *Image Vis Comput.* 2001;19:1023-1040.
- 425 22. Montagnat J, Delingette H. 4D deformable models with temporal constraints: application to 4D cardiac image segmentation. *Med Image Anal* 2005;9:87-100.
23. Schreibmann E, Xing L. Image registration with auto-mapped control volumes. *Med Phys* 2006;33:1165-1179.
24. Ibanez L, Schroeder W, Ng L. ITK Software Guide: Kitware Inc.; 2003.
- 430 25. Schroeder W, Martin K, Lorensen B. The Visualization Toolkit: An Objective-Oriented Approach To 3D Graphics.
26. Broyden CG. The convergence of a class of double-rank minimization algorithms 2, the new algorithm. *J of Inte Inst For Math And Applications* 1970;6:222-231.
27. Fletcher R. A new approach to variable-metric algorithms. *Computer Journal* 1970;13:317-322.
- 435 28. Goldfarb D. A family of variable-metric algorithms derived by variational means. *Math of Computation* 1970;24:23-26.
29. Shanno DF. Conditioning of Quasi-Newton methods for function minimization. *Math of Computation* 1970;24:647-656.
- 440 30. Liu DC, Nocedal J. On the limited memory BFGS method for large scale optimization. *Math Program* 1989;45:503-528.
31. Schreibmann E, Xing L. Narrow band deformable registration of prostate magnetic resonance imaging, magnetic resonance spectroscopic imaging, and computed tomography studies. *Int J Radiat Oncol Biol Phys* 2005;62:595-605.
- 445 32. Rietzel E, Chen GT, Choi NC, Willet CG. Four-dimensional image-based treatment planning: Target volume segmentation and dose calculation in the presence of respiratory motion. *Int J Radiat Oncol Biol Phys* 2005;61:1535-1550.
33. Lu W, Olivera GH, Chen Q, Chen M, Ruchala KJ. Automatic re-contouring in 4D radiotherapy. *Phys Med Biol* 2006;51:1077-1099.

450

455

460

465

470

FIGURE CAPTIONS

Figure 1. A transverse view of the template image of patient 1 with the manually drawn contour of the lung (green curve) and the control volumes (pink squares). Each control volume is typically ~1cm in cubic shape in our study.

475

Figure 2. Flow chart of the automatic contour mapping process for 4D CT images.

Figure 3. Depiction of algorithm validation using a simulated deformation matrix on the image and the contour. (a) Original image and contour (green); (b) Artificially deformed image and contour (yellow) and the mapped contour using the novel algorithm (red). The blue contour is the contour result after rigid mapping. The original contour in green is overlaid on (b) for reference purpose.

480

Figure 4. (a) Energies (self, interaction and total energy) as a function of the searching
485 points in the surrounding area of a randomly selected control volume in the target image
(see text for details). (b) The metric value versus step during the course of the contour
mapping. The metric values are the average NCC of all the involved control volumes.

Figure 5. The manually delineated contours for three lung cancer patients on 3D CT
490 images. (a) patient 1 and (b) patient 2 and (c) patient 3. In (a) and (b) the red curves
represent the body contour and the green represent the lungs. In (c) the body contour is
shown in red and the lung in yellow. Purple and green curves represent PTV and GTV
contours, respectively.

495 Figure 6. A transverse view of CT images and contours for the first patient. The top row
(a), (b) and (c) are for phase 4; the middle row (d), (e) and (f) for phase 7, and the bottom
row (g), (h) and (i) for phase 10. For each phase, we show different slices. For example,
(a), (d) and (g) are slices at the position of 190 cm; (b), (e) and (h) are slices at 240 cm
500 and (c), (f) and (i) at 300 cm. The contours in green were the manually delineated
contours (template contours) and were superimposed in the images for reference purposes.
The blue contours are the rigidly mapped contours. The contours in red are the final
results (target contours).

505 Figure 7. A transverse view of CT images and contours for the second patient. The top
row (a), (b) and (c) are for phase 2; the middle row (d), (e) and (f) for phase 5, and the
bottom row (g), (h) and (i) for phase 8. For each phase, we show different slices. For
example, (a), (d) and (g) are slices at the position of 140 cm; (b), (e) and (h) are slices at
175 cm and (c), (f) and (i) at 250 cm. The contours in green were the manually delineated
contours (template contours) and were superimposed in the images for reference purposes.
510 The blue contours are the rigidly mapped contours. The contours in red are the final
results (target contours).

Figure 8. A transverse view of the CT images and GTV contours for patient 3 on the
same slice at different phases. (a), (b), (c), and (d) correspond to phases 2, 4, 7, and 9,

515 respectively. The manually drawn contours are shown in green, rigidly mapped contours
in blue, and the target contours in red.

Figure 1

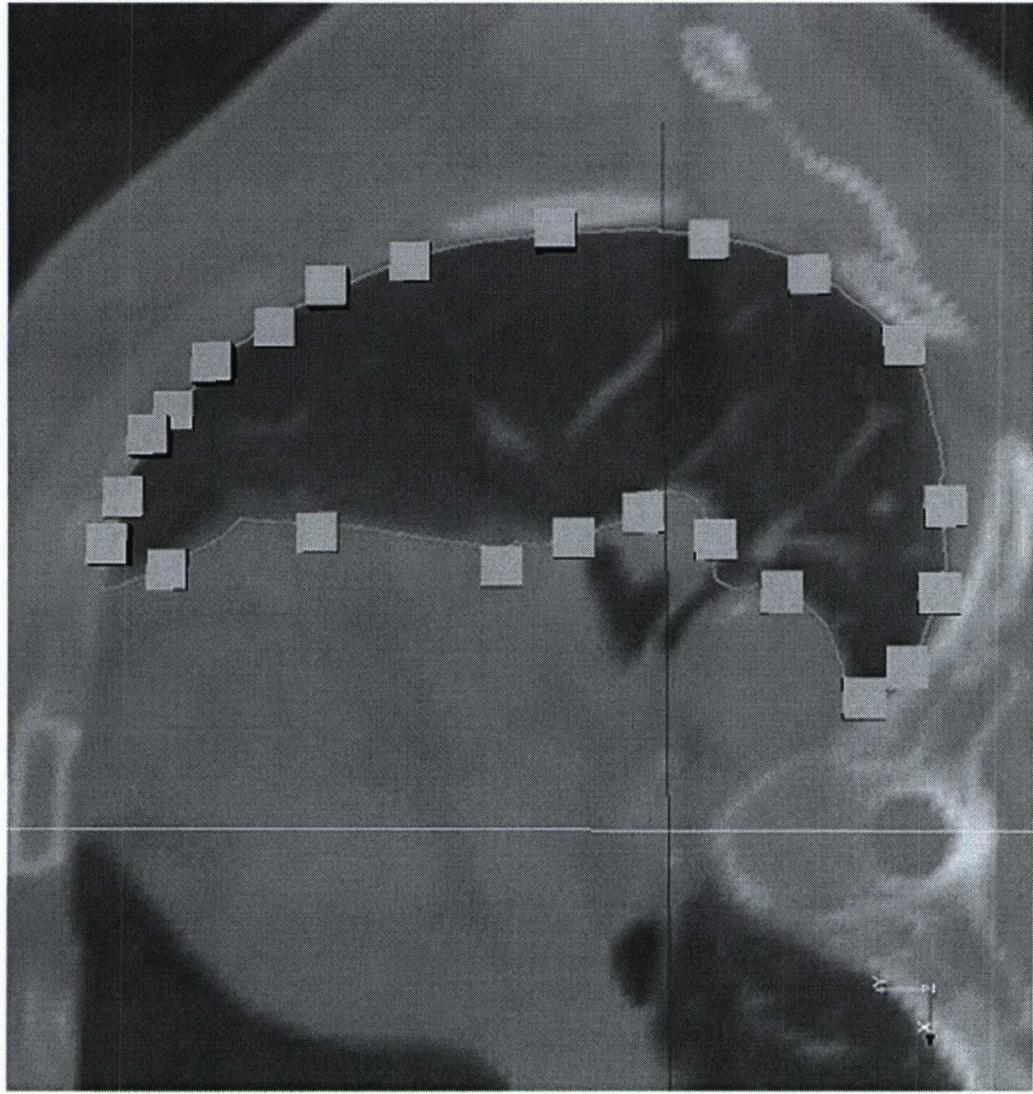


Figure 2

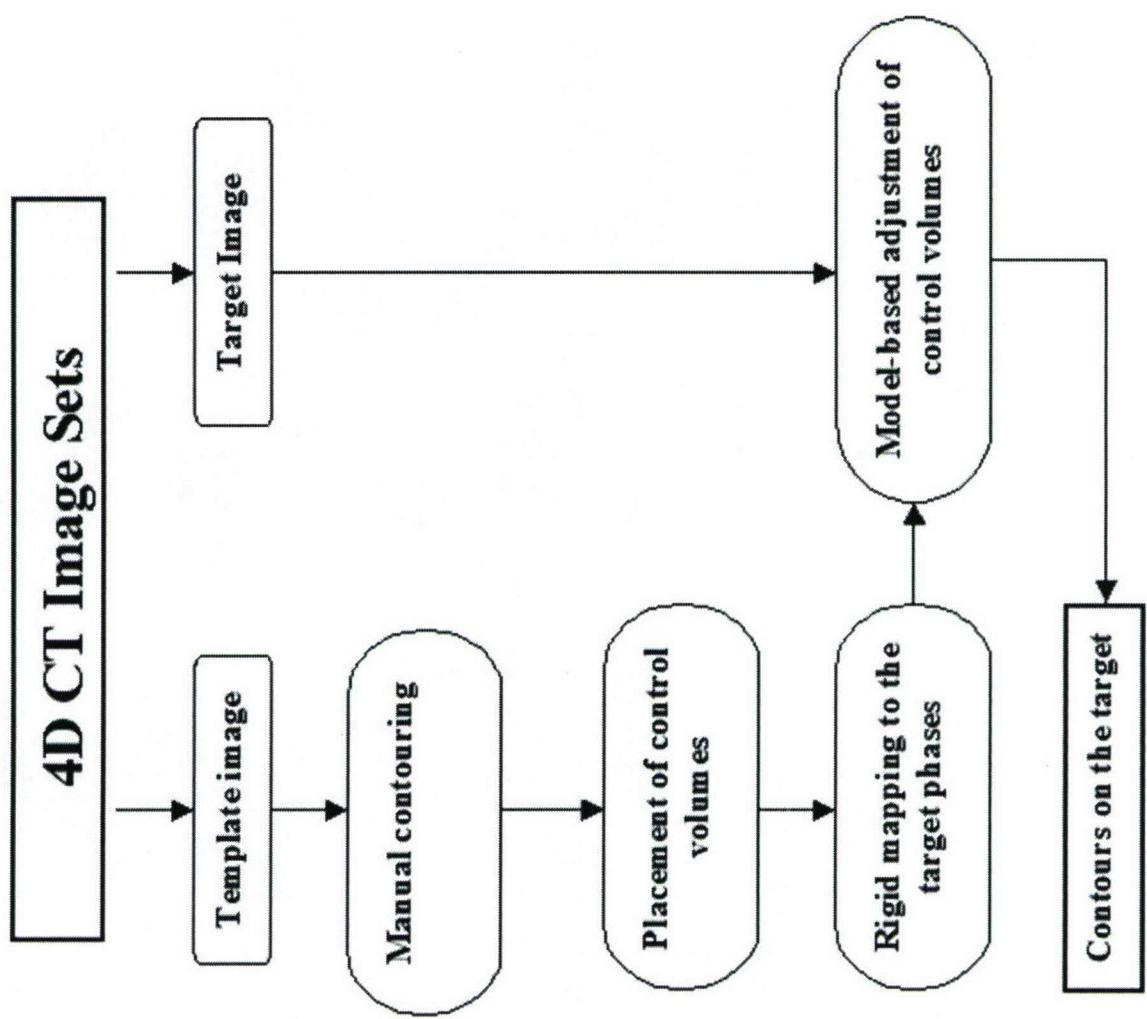


Figure 3

(b)



(a)

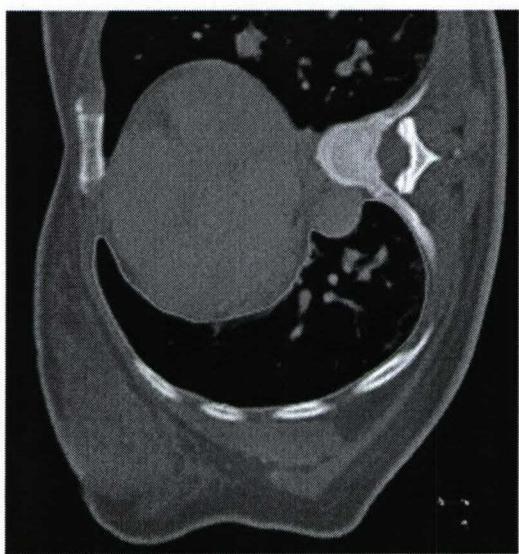


Figure 4

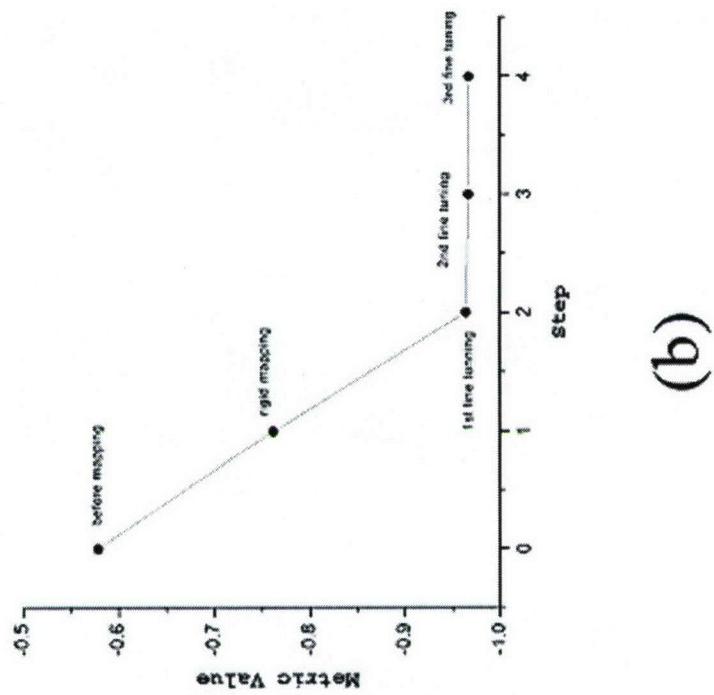
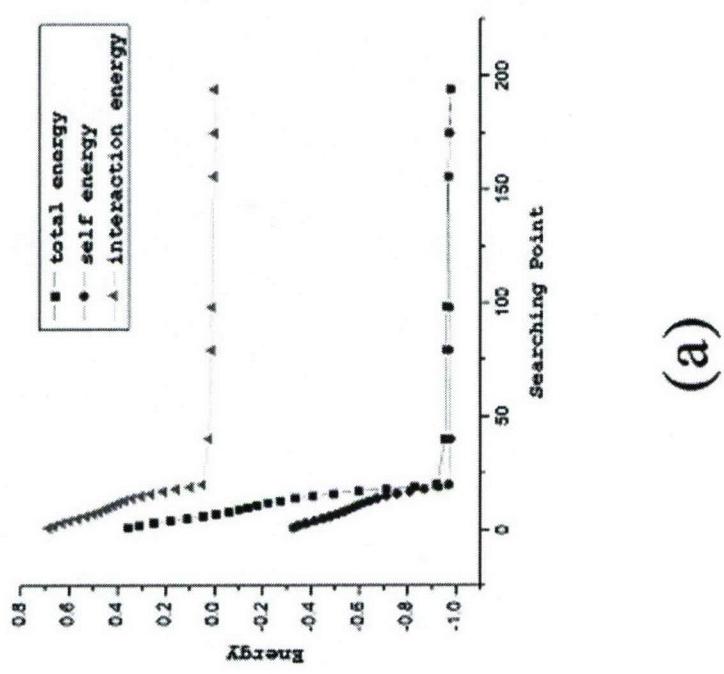


Figure 5

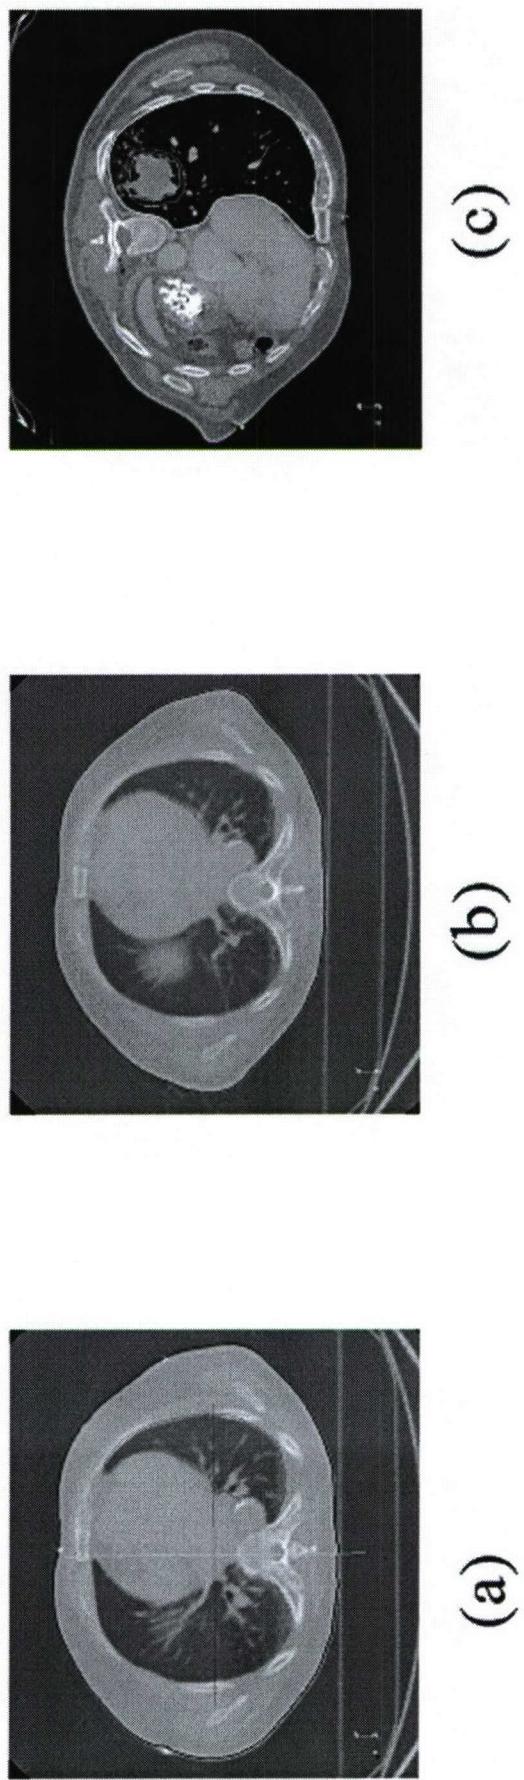


Figure 6

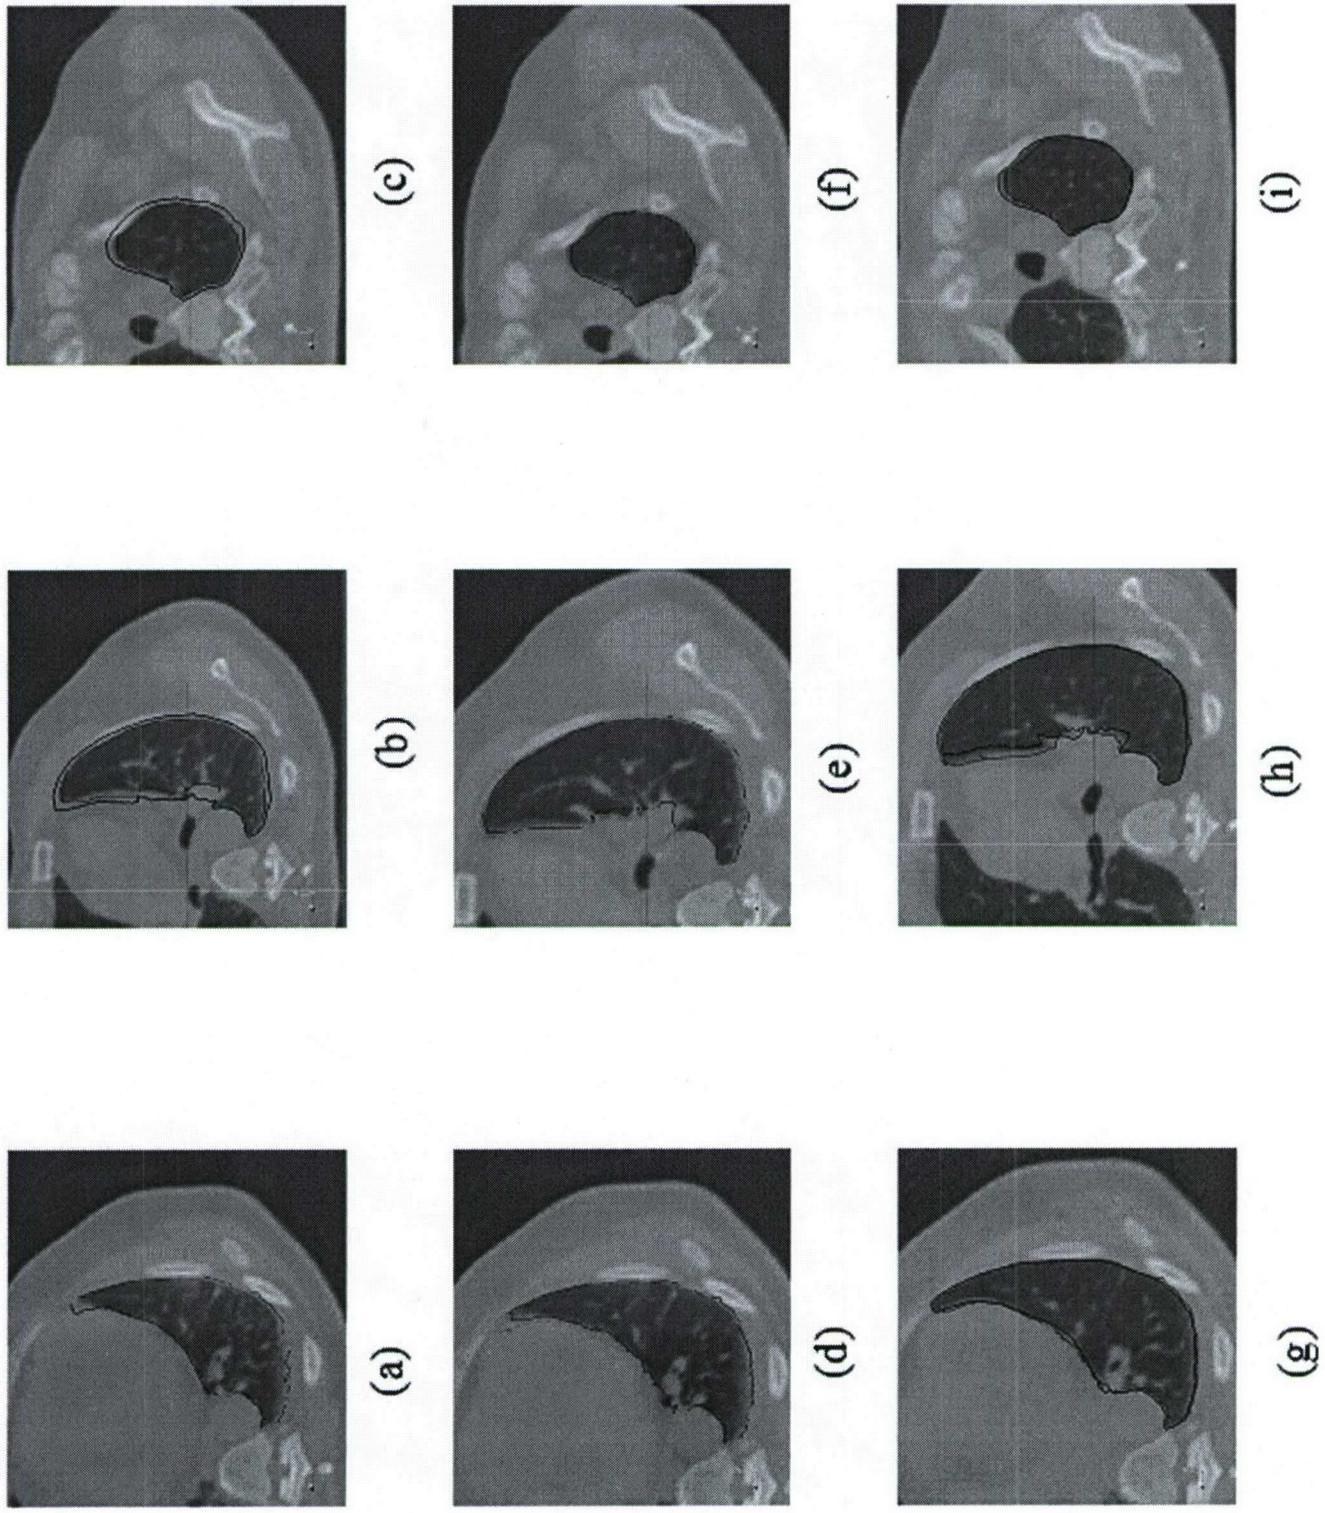


Figure 7

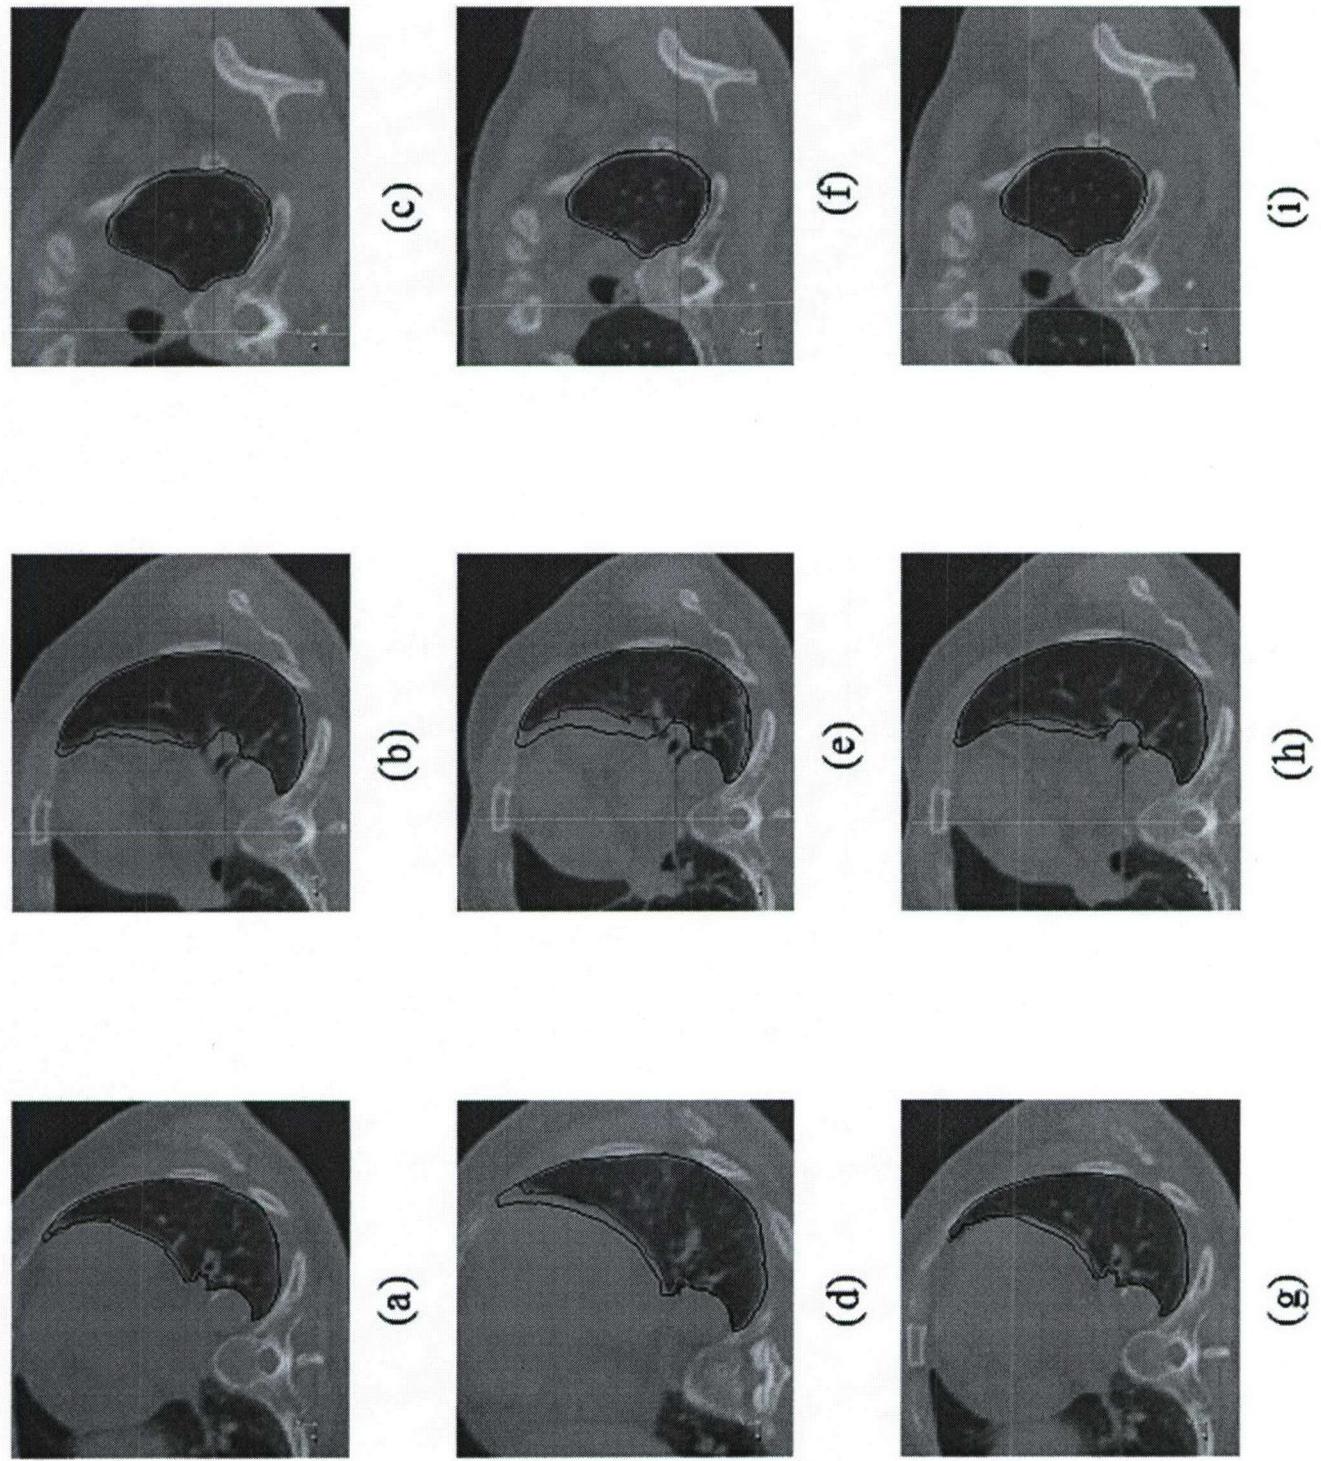
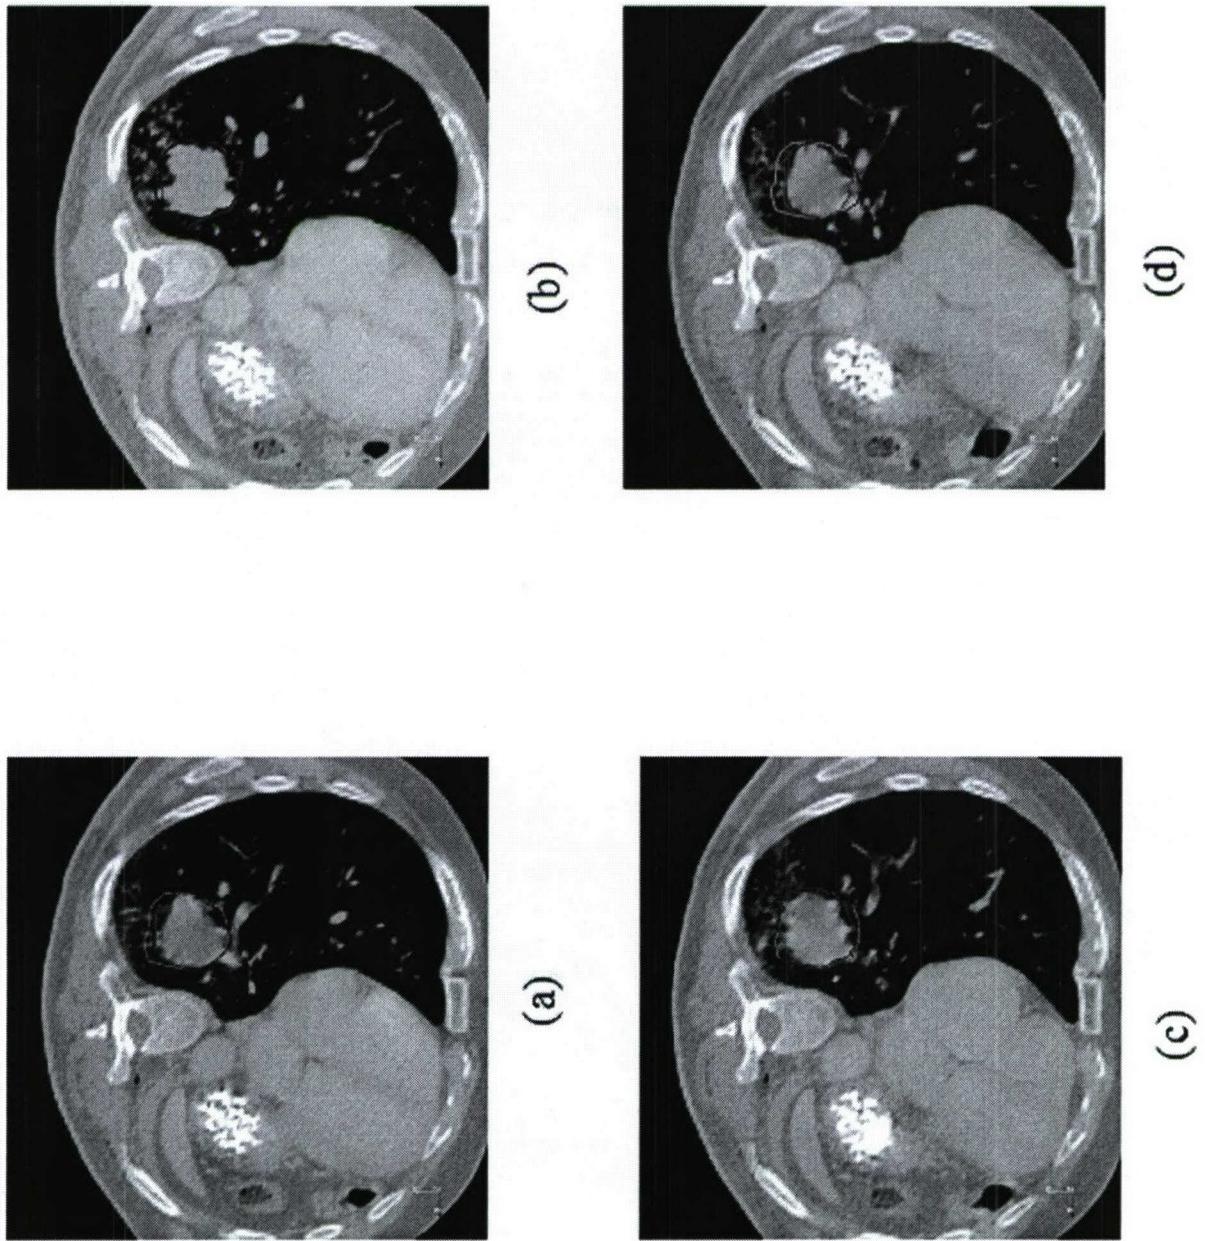


Figure 8



Multiscale Deformable Registration

Dana Paquin*, Doron Levy, and Lei Xing

Abstract— Multiscale image registration techniques are presented for the deformable registration of medical images. These techniques are shown to be particularly effective for registration problems in which one or both of the images to be registered contains significant levels of noise. Experiments using B-splines deformable registration models demonstrate that ordinary deformable registration techniques fail to produce accurate results in the presence of significant levels of noise. Further experiments in which classical denoising algorithms are applied prior to registration show that registration fails in this case for high levels of noise, as well. The hierarchical multiscale image decomposition of E. Tadmor, S. Nezzar, and L. Vese, *A multiscale image representation using hierarchical (BV, L^2) decompositions*, *Multiscale Modeling and Simulations*, vol. 2, no. 4, pp. 554–579, 2004, is presented, and multiscale image registration algorithms are developed based on the multiscale decomposition. An accurate registration of noisy images is achieved by obtaining a hierarchical multiscale decomposition of the images and registering the resulting components. This approach enables a successful registration of images that contain noise well beyond the level at which ordinary deformable registration fails. Further, an iterated multiscale registration algorithm is presented which improves the accuracy and computational efficiency of other registration methods. The accuracy and efficiency of the multiscale registration techniques are demonstrated using both digital phantom and clinical case studies in two and three dimensions.

Index Terms— Image registration, deformable registration, noise, multiscale decomposition.

I. INTRODUCTION

IMAGE registration is the process of determining the optimal spatial transformation that maps one image to another. Image registration is necessary, for example, when images of the same object are taken at different times, from different imaging devices, or from different perspectives. The two images to be registered, called the fixed and moving images, are the input to the registration algorithm,

and the output is the optimal transformation that maps the moving image to the fixed image. Applications of image registration include image-guided radiation therapy (IGRT), image-guided surgery, functional MRI analysis, and tumor detection, as well as many non-medical applications, such as computer vision, pattern recognition, and remotely sensed data processing (see [4] and the references therein).

Image registration models are classified into two main categories according to the transformation type: affine and deformable. Affine image registration models assume that the transformation that maps the moving image to the fixed image consists only of translations and rotations, while deformable models allow localized stretching of images. While affine models are sufficient in certain circumstances, many registration problems, particularly in medical applications, are non-rigid. For example, respiratory motion causes non-rigid, or deformable, distortion of the lungs and other organs. As another example, image-guided neurosurgery procedures require deformable registration of pre- and intra-operative images of the brain [14], [20]. For additional applications of deformable registration, see [11], [17], [18] and the references therein.

This paper extends our previous work [13], a multiscale approach to rigid registration, to deformable registration problems. The structure of this paper is as follows. In Section II, we provide a brief overview of the image registration problem and discuss deformable registration techniques. In Section III, we describe the problem of deformable image registration in the presence of noise, and briefly discuss standard denoising techniques. In Section IV, we review the hierarchical multiscale image decomposition of [19], and we present three multiscale image registration algorithms based on the decomposition. In Section V, we demonstrate the failure of ordinary deformable registration techniques when one or more of the images to be registered contains noise, as well as the failure of ordinary techniques even when the noisy images are denoised using classical filtering methods prior to registration. The accuracy and efficiency of our multiscale registration techniques are presented in Section V with several image registration experi-

*Asterisk indicates corresponding author.

D. Paquin is with the Department of Mathematics, Stanford University, Stanford, CA 94305-2125 (e-mail: paquin@math.stanford.edu).

D. Levy is with the Department of Mathematics, Stanford University, Stanford, CA 94305-2125 (e-mail: dlevy@math.stanford.edu).

L. Xing is with the Department of Radiation Oncology, Stanford University, Stanford, CA 94305-5847 (e-mail: lei@reyes.stanford.edu).

ments in both two and three dimensions. Concluding remarks are given in Section VI.

II. MATHEMATICAL FORMULATION OF THE REGISTRATION PROBLEM

Given a *fixed* and a *moving* image, the registration problem is the process of finding an *optimal transformation* that brings the moving image into spatial alignment with the fixed image. While this problem is easy to state, it is difficult to solve. The main source of difficulty is that the problem is ill-posed, which means, for example, that the problem may not have a unique solution. For a detailed overview of the image registration problem and various image registration techniques, see [12].

To formulate the registration problem mathematically, a two-dimensional gray-scale image f is a mapping which assigns to every point $x \in \Omega \subset \mathbb{R}^2$ a gray value $f(x)$ (called the intensity value of the image at the point x). We will consider images as elements of the space $L^2(\mathbb{R}^2)$. Any registration algorithm has three main components:

- 1) The *transformation model* which specifies the way in which the moving image can be transformed to correspond to the fixed image.
- 2) The *distance measure*, or metric, used to compare the fixed and moving images.
- 3) The *optimization process* that varies the parameters of the transformation model in such a way that the transformation produced by the registration process is optimal.

Given a distance measure $D : (L^2(\mathbb{R}^2))^2 \rightarrow \mathbb{R}$ and two images $f(x), m(x) \in L^2(\mathbb{R}^2)$, the solution ϕ of the registration problem is given by the following minimization problem:

$$\phi = \underset{\psi: \mathbb{R}^2 \rightarrow \mathbb{R}^2}{\operatorname{argmin}} D(f(x), m(\psi(x))), \quad (1)$$

where ψ is in the specified space of transformation models. Examples of commonly used distance measures are mean squares, normalized correlation, and mutual information. Examples of typical transformation models are rigid, affine, polynomial, and spline transformations [12]. To minimize $D(f, m(\psi))$, we must choose an optimizer which controls the minimization. The most commonly used optimization techniques in image registration are gradient descent and regular step gradient descent methods. The implementation of the registration algorithm works in the following way: at each iteration, the distance D between the two images is computed. The specified transformation is then applied to the moving image, and the distance

between the images is recomputed. In theory, this process continues until the distance is minimized (or maximized in certain cases), though in practice a stopping criterion is often applied.

Much of the current work in medical image registration is focused on the deformable case. Although deformable image registration clearly allows for more flexibility in the types of images and applications in which it can be used than rigid registration, deformable registration techniques require significantly more computation time than rigid registration techniques, and involve the determination of a very large number of parameters. In this paper, we shall focus on the problem of *deformable registration in the presence of noise*.

A. Deformable registration techniques

Spline-based transformation models are among the most common and important transformation models used in non-rigid registration problems [6]. Spline-based registration algorithms use *control points* in the fixed and moving images and a spline function to define transformations away from these points. The two main spline models used in registration are *thin-plate splines* and *B-splines*. Thin-plate splines have the property that each control point has a global influence on the transformation. This can be a disadvantage because it limits the ability of the transformation model to model localized deformations. In addition, the computation time required for a thin-plate spline-based registration algorithm increases significantly as the number of control points increases. See [3] for an overview of thin-plate splines.

In contrast, B-splines are only defined in the neighborhood of each control point. Thus perturbing the position of one control point affects the transformation only in a neighborhood of that point. As a result, B-spline-based registration techniques are more computationally efficient than thin-plate splines, especially for a large number of control points. See [9] and [10] for a detailed description of B-splines transformation models. In this paper, we shall use deformable registration algorithms based on B-spline deformation models. To define the spline-based deformation model, let $\Omega = \{(x, y) \mid 0 \leq x \leq X, 0 \leq y \leq Y\}$ denote the domain of the image volume. Let α denote a $n_x \times n_y$ mesh of control points $\alpha_{i,j}$ with uniform spacing δ . Then the B-spline deformation model can be written as the 2-D tensor product of 1-D cubic B-splines:

$$\phi(x, y) = \sum_{l=0}^2 \sum_{m=0}^2 B_l(x) B_m(y) \alpha_{i+l, j+m}, \quad (2)$$

where $i = \lfloor x/n_x \rfloor - 1$, $j = \lfloor y/n_y \rfloor - 1$, and B_l represents the l -th basis of the B-spline:

$$\begin{aligned} B_0(u) &= \frac{1}{6}(1-u)^3, \\ B_1(u) &= \frac{1}{6}(3u^3 - 6u^2 + 4), \\ B_2(u) &= \frac{1}{6}(-3u^3 + 3u^2 + 3u + 1). \end{aligned}$$

Changing the control point $\alpha_{i,j}$ affects the transformation only in a local neighborhood of $\alpha_{i,j}$. The control points α act as parameters of the B-spline deformation model, and the degree of non-rigid deformation which can be modeled depends on the resolution of the mesh of control points α . A large spacing of control points allows modeling of global non-rigid deformation, while a small spacing of control points allows modeling of local non-rigid deformations. Additionally, the number of control points determines the number of degrees of freedom of the transformation model, and hence, the computational complexity.

In Section V, we demonstrate the accuracy of B-splines deformable registration techniques for the registration of non-noisy images. Finally, we note in passing that there are additional deformable registration techniques such as elastic models [2], viscous fluid models [5], and finite element models [7].

III. REGISTRATION IN THE PRESENCE OF NOISE

A. Speckle noise model

In imaging, the term noise refers to random fluctuations in intensity values that occur during image capture, transmission, or processing, and that may distort the information given by the image. Image noise is not part of the ideal signal and may be caused by a wide range of sources, such as detector sensitivity, environmental radiation, transmission errors, discretization effects, etc. The specific type of noise present in a given image depends on the imaging modality and even on the specific scanner. Common types of noise present in medical images are impulse ('salt & pepper'), speckle, Gaussian, and Rician (present in MR images). In this paper, we will study the problem of image registration in the presence of high levels of *speckle* noise.

Although we present results here only for this specific type of noise, we would like to note that the algorithm proposed in this work is not noise specific, and is extendable to register any types of noisy images. In particular, we have conducted numerous registration simulations with other noise models and have obtained results similar to those presented here. See, for example, our results on rigid registration [13].

Speckle noise, or multiplicative noise, is a type of noise that occurs commonly in medical imaging. In particular, speckle noise is often found in ultrasound images [1]. It is defined by the following model. We let $s(x)$ denote the actual image, and $f(x)$ the observed image. Then

$$f(x) = s(x) + \eta(0, \delta) \cdot s(x), \quad (3)$$

where $\eta(0, \delta)$ is uniformly distributed random noise of mean 0 and variance δ . We add speckle noise of increasing variance to the image S , as illustrated in Figure 4. For a given noise variance δ , we denote the noisy image S_δ . In Section V, we present the registration results obtained upon registering the noisy images S_δ with the original image I for increasing values of δ .

B. Classical denoising techniques

Since we are considering the problem of image registration in the presence of noise, it is natural to consider whether or not the application of standard denoising algorithms prior to registration enables a successful deformable registration of noisy images. Image denoising is a fundamental problem in image processing, and there has been much research and progress on the subject. As our primary interest is the problem of image registration of noisy images, and not denoising, we do not focus on the general problem of image denoising, but instead consider a few of the most common and computationally simple denoising techniques.

Spatial filtering is the traditional approach to removing noise from images. Spatial filters use the assumption that noise occupies the higher regions of the frequency spectrum, and thus attenuate high spatial frequencies. Local spatial filtering is a process in which the value of a given pixel in the filtered image is computed by applying some algorithm to the pixel values in a neighborhood of the given pixel. Typical implementations of spatial filters include mean filtering, median filtering, and Gaussian smoothing. See [21] for an overview of classical spatial filtering techniques.

Mean filtering computes the value of each output pixel by computing the statistical mean of the neighborhood of the corresponding input pixel. That is, the intensity value of each pixel is replaced with the mean, or average, value of its neighbors (including itself). The neighborhood size is typically taken to be a 3×3 or a 5×5 square grid (or kernel); larger neighborhoods result in more severe smoothing. In this paper, we use a 3×3 square kernel to perform mean filtering. Applying a mean filter to a noisy image reduces the amount of variation in gray-level intensity between pixels. Although this filter is computationally easy to implement, it is sensitive to the presence of outliers.

Median filtering computes the value of each output pixel by computing the statistical median of the neighborhood of the corresponding input pixel. The median filter has two main advantages over the mean filter. First, the median is a more robust average than the mean, and thus is less sensitive to a single very unrepresentative pixel in a neighborhood. Second, since the median value must actually be the value of one of the pixels in the neighborhood, the median filter does not create unrealistic pixel values when the filter straddles an edge. For this reason, the median filter is much more effective at preserving sharp edges than the mean filter. As with the mean filter, we use a 3×3 square kernel to perform median filtering.

Convolution with a Gaussian kernel is another commonly used spatial filtering technique. The 2-dimensional Gaussian kernel has the form

$$G(x, y) = \frac{1}{2\sigma^2} e^{-\frac{x^2+y^2}{2\sigma^2}}, \quad (4)$$

where σ is the standard deviation of the distribution, which we assume has mean zero. The Gaussian filter computes a weighted average of each pixel's neighborhood, with the average weighted toward the value of the central pixels, in contrast to the mean filter's uniformly weighted average. In Section V, we present the registration results obtained upon applying classical denoising techniques to noisy images prior to registration.

IV. MULTISCALE REGISTRATION ALGORITHMS

A. The multiscale decomposition

The multiscale registration techniques to be discussed in this paper are based on the multiscale image representation using the hierarchical (BV, L^2) decompositions of [19]. This multiscale decomposition will provide a hierarchical expansion of an image that separates the essential features of the image (such as large shapes and edges) from

the fine scales of the image (such as details and noise). The decomposition is hierarchical in the sense that it will produce a series of expansions of the image that resolve increasingly finer scales, and hence include increasing levels of detail. We will eventually apply the multiscale decomposition algorithm to the problem of image registration in the presence of noise, and will demonstrate the accuracy of the multiscale registration technique for noisy images.

We will use the following mathematical spaces in the decomposition algorithm. The space of functions of bounded variation, BV , is defined by:

$$BV = \{f \mid \|f\|_{BV} := \sup_{h \neq 0} |h|^{-1} \|f(\cdot + h) - f(\cdot)\|_{L^1} < \infty\}.$$

We will also use the Sobolev space $W^{-1,\infty}$ with norm given by:

$$\|f\|_{W^{-1,\infty}} := \sup_g \left[\frac{\int f(x)g(x) dx}{\|g\|_{W^{1,1}}} \right],$$

where $\|g\|_{W^{1,1}} := \|\nabla g\|_{L^1}$.

Define the J -functional $J(f, \lambda)$ as follows:

$$J(f, \lambda) := \inf_{u+v=f} \lambda \|v\|_{L^2}^2 + \|u\|_{BV}, \quad (5)$$

where $\lambda > 0$ is a scaling parameter that separates the L^2 and BV terms. This functional $J(f, \lambda)$ was introduced in the context of image processing by Rudin, Osher, and Fatemi [15]. Let $[u_\lambda, v_\lambda]$ denote the minimizer of $J(f, \lambda)$. The BV component, u_λ , captures the coarse features of the image f , while the L^2 component, v_λ , captures the finer features of f such as noise. This model is effective in denoising images while preserving edges, though it requires prior knowledge on the noise scaling λ .

Tadmor, *et al.* proposed in [19] an alternative point of view in which the minimization of $J(f, \lambda)$ is interpreted as a decomposition $f = u_\lambda + v_\lambda$, where u_λ extracts the edges of f and v_λ extracts the textures of f . This interpretation depends on the scale λ , since texture at scale λ consists of edges when viewed under a refined scale. Upon decomposing $f = u_\lambda + v_\lambda$, we proceed to decompose v_λ as follows:

$$v_\lambda = u_{2\lambda} + v_{2\lambda},$$

where

$$[u_{2\lambda}, v_{2\lambda}] = \operatorname{arginf}_{u+v=v_\lambda} J(v_\lambda, 2\lambda).$$

Thus we obtain a two-scale representation of f given by $f \cong u_\lambda + u_{2\lambda}$. Repeating this process results in the following hierarchical multiscale decomposition of f . Starting with an initial scale $\lambda = \lambda_0$, we obtain an initial decomposition of the image f :

$$f = u_0 + v_0,$$

$$[u_0, v_0] = \underset{u+v=f}{\operatorname{arginf}} J(f, \lambda_0).$$

We then refine this decomposition to obtain

$$v_j = u_{j+1} + v_{j+1},$$

$$[u_{j+1}, v_{j+1}] = \underset{u+v=v_j}{\operatorname{arginf}} J(v_j, \lambda_0 2^{j+1}), \quad j = 0, 1, \dots$$

After k steps of this process, we have:

$$\begin{aligned} f &= u_0 + v_0 \\ &= u_0 + u_1 + v_1 = \dots \\ &= u_0 + u_1 + \dots + u_k + v_k, \end{aligned} \tag{6}$$

which is a multiscale image decomposition $f \sim u_0 + u_1 + \dots + u_k$, with a residual v_k . As k increases, the u_k components resolve edges with increasing scales $\lambda_k = \lambda_0 2^k$.

1) Implementation of the multi-scale decomposition: As described in [19], the initial scale λ_0 should capture the smallest oscillatory scale in f , given by

$$\frac{1}{2\lambda_0} \leq \|f\|_{W^{-1,\infty}} \leq \frac{1}{\lambda_0}. \tag{7}$$

However, in practice, we may not be able to determine the size of $\|f\|_{W^{-1,\infty}}$, so we determine the initial choice of λ_0 experimentally. Experimentally, the initial scale λ_0 should be chosen such that the decomposition produces a hierarchy of scales containing increasing levels of detail. If the initial λ_0 is too small, then the minimizer of $J(f, \lambda_0)$ will just be the trivial pair $[0, f]$, i.e. the first scale decomposition will simply be the original image f . In this case, λ_0 should be increased (typically by dyadic multiples) until the energy estimate given in Eq. (7) holds. Following [19], for the applications presented in this paper, we will use $\lambda_0 = 0.01$ and $\lambda_j = \lambda_0 2^j$. In our simulations, we have found that $\lambda_0 = 0.01$ works well for most medical images (such as those used in this paper).

We follow the numerical algorithm of [19] for the construction of our hierarchical decomposition. In each step, we use finite-difference discretization of the Euler-Lagrange equations associated with the $J(v_j, \lambda_{j+1})$ to obtain the next term, u_{j+1} , in the decomposition of the image f . Due to the singularity when $|\nabla u_\lambda| = 0$, we replace $J(f, \lambda)$ by the regularized functional $J^\epsilon(f, \lambda) :=$

$$\inf_{u+v=f} \{\lambda \|v\|_{L^2}^2 + \int_{\Omega} \sqrt{\epsilon^2 + |\nabla u|^2} dx dy\}, \tag{8}$$

and at each step, we find the minimizer u_λ of J^ϵ . The Euler-Lagrange equation for $J^\epsilon(f, \lambda)$ is

$$u_\lambda - \frac{1}{2\lambda} \operatorname{div} \left(\frac{\nabla u_\lambda}{\sqrt{\epsilon^2 + |\nabla u_\lambda|^2}} \right) = f \quad \text{in } \Omega,$$

with the Neumann boundary conditions:

$$\frac{\partial u_\lambda}{\partial n} \Big|_{\partial\Omega} = 0, \tag{9}$$

where $\partial\Omega$ is the boundary of the domain Ω and n is the unit outward normal. We thus obtain an expansion $f \sim \sum_{j=0}^k u_j$, where the u_j are constructed as approximate solutions of the recursive relation given by the following elliptic PDE:

$$\begin{aligned} u_{j+1} - \frac{1}{2\lambda_{j+1}} \operatorname{div} \left(\frac{\nabla u_{j+1}}{\sqrt{\epsilon^2 + |\nabla u_{j+1}|^2}} \right) \\ = -\frac{1}{2\lambda_j} \operatorname{div} \left(\frac{\nabla u_j}{\sqrt{\epsilon^2 + |\nabla u_j|^2}} \right). \end{aligned} \tag{10}$$

To numerically implement the method, we cover the domain Ω with a grid of cell size h ($x_i := ih, y_j := jh$), and discretize the elliptic PDE of Eq. (10) as follows::

$$\begin{aligned} u_{i,j} &= f_{i,j} \\ &+ \frac{1}{2\lambda h^2} \left[\frac{u_{i+1,j} - u_{i,j}}{\sqrt{\epsilon^2 + (D_{+x} u_{i,j})^2 + (D_{0y} u_{i,j})^2}} \right] \\ &- \frac{1}{2\lambda h^2} \left[\frac{u_{i,j} - u_{i-1,j}}{\sqrt{\epsilon^2 + (D_{-x} u_{i,j})^2 + (D_{0y} u_{i-1,j})^2}} \right] \\ &+ \frac{1}{2\lambda h^2} \left[\frac{u_{i,j+1} - u_{i,j}}{\sqrt{\epsilon^2 + (D_{0x} u_{i,j})^2 + (D_{+y} u_{i,j})^2}} \right] \\ &- \frac{1}{2\lambda h^2} \left[\frac{u_{i,j} - u_{i,j-1}}{\sqrt{\epsilon^2 + (D_{0x} u_{i,j-1})^2 + (D_{-y} u_{i,j})^2}} \right], \end{aligned} \tag{11}$$

where D_+ , D_- , and D_0 denote the forward, backward, and centered divided differences, respectively. To solve the discrete regularized Euler-Lagrange equations (11), we use the Gauss-Siedel iterative method to obtain:

$$\begin{aligned} u_{i,j}^{n+1} &= f_{i,j} \quad (12) \\ &+ \frac{1}{2\lambda h^2} \left[\frac{u_{i+1,j}^n - u_{i,j}^{n+1}}{\sqrt{\epsilon^2 + (D_{+x} u_{i,j}^n)^2 + (D_{0y} u_{i,j}^n)^2}} \right] \\ &- \frac{1}{2\lambda h^2} \left[\frac{u_{i,j}^{n+1} - u_{i-1,j}^n}{\sqrt{\epsilon^2 + (D_{-x} u_{i,j}^n)^2 + (D_{0y} u_{i-1,j}^n)^2}} \right] \\ &+ \frac{1}{2\lambda h^2} \left[\frac{u_{i,j+1}^n - u_{i,j}^{n+1}}{\sqrt{\epsilon^2 + (D_{0x} u_{i,j}^n)^2 + (D_{+y} u_{i,j}^n)^2}} \right] \\ &- \frac{1}{2\lambda h^2} \left[\frac{u_{i,j}^{n+1} - u_{i,j-1}^n}{\sqrt{\epsilon^2 + (D_{0x} u_{i,j-1}^n)^2 + (D_{-y} u_{i,j}^n)^2}} \right] \end{aligned}$$

To satisfy the Neumann boundary conditions (9), we first reflect f outside Ω by adding grid lines on all sides of Ω . As the initial condition, we set $u_{i,j}^0 = f_{i,j}$. We iterate this numerical scheme for $n = 0, 1, \dots, N$ until $\|u^{n_\infty} - u^{n_\infty-1}\|$ is less than some preassigned value so that $u_{i,j}^{n_\infty}$ is an accurate approximation of the fixed point steady solution u_λ .

Finally, we denote the final solution $u_\lambda := \{u_{i,j}^{n_\infty}\}_{i,j}$. To obtain the hierarchical multiscale decomposition, we reiterate this process, each time updating f and λ in the following way:

$$\begin{aligned} f_{\text{new}} &\leftarrow f_{\text{current}} - u_\lambda, \\ \lambda_{\text{new}} &\leftarrow 2\lambda_{\text{current}}. \end{aligned} \quad (13)$$

Taking $\lambda_j = \lambda_0 2^j$, we obtain after k steps a hierarchical multiscale decomposition $f = u_{\lambda_0} + u_{\lambda_1} + \dots + u_{\lambda_k} + v_{\lambda_k}$, where we write $u_{\lambda_j} = u_j$. We call the u_j , $j = 1, 2, \dots, k$ the components of f , and the v_k the residuals. For ease of notation, given an image f , we let $C_k(f)$ denote the k^{th} scale of the image f , $k = 1, \dots, m$:

$$C_k(f) = \sum_{i=0}^{k-1} u_k(f). \quad (14)$$

Thus $C_k(A)$ will denote the k^{th} scale of the image A , and $C_k(B)$ will denote the k^{th} scale of image B .

There are several stopping criteria which can be used to determine the number of hierarchical steps m to use. One is to stop the decomposition when

the amount of texture $\|v_{2k-1}\|_W^{-1,\infty}$ is below a certain value. Another is to stop when the energy $\|u_k - u_{k-1}\|_{L^2\Omega}$ is below a certain value. For the medical images that we have considered, we have found that using $m = 8$ hierarchical steps produces a multiscale decomposition in which $C_1(f)$ contains only very coarse scales and $C_8(f)$ is close to the original image f .

As an example, we apply this multiscale decomposition to the noisy deformed image $S_{0.6}$, using 8 hierarchical steps. The results are shown in Figure 7.

B. Multiscale registration

In this section, we present multiscale image registration techniques that are based on the hierarchical multiscale decomposition of [19] reviewed in Section IV-A. For the general setup, consider two images A (the fixed image) and B (the moving image), and suppose that we want to register image B with image A . Suppose that one or both of the images contains a significant amount of noise. If only one of the images is noisy, we assume that it is image B . For each of the three algorithms presented below, we first apply the multiscale hierarchical decomposition to both images. Let m denote the number of hierarchical steps used for the multiscale decompositions. We presented Algorithms I and II in the context of rigid registration in [13].

1) Algorithm I: Single-node registration: In our single-node multiscale registration algorithm, Algorithm I, we register the k^{th} scale $C_k(B)$ of image B with the image A , for $k = 1, \dots, m$. We refer to this algorithm as a *one-node multiscale registration algorithm* because in each of the m registrations prescribed by the algorithm, the moving image is always the image A . That is, in the single-node algorithm, we use only the multiscale components of the fixed image B . Since this algorithm considers scales only of the noisy (moving) image, we expect that it will be particularly successful when only one of the images to be registered is noisy.

Recall that $C_1(B)$ contains only the coarsest scales of the image B , and as k increases, $C_k(B)$ contains increasing levels of detail (and hence, noise) of the image B . Thus, we expect that the registration of $C_k(B)$ with A should be more accurate than ordinary registration for the first few values of k . As k increases, however, we expect that eventually the scale $C_k(B)$ will become too noisy to register successfully.

Upon registering $C_k(B)$ with A for $k = 1, 2, \dots, m$, we can estimate the actual transformation that maps the images to one another either by

computing a weighted average of the transformations produced by these m registrations, or simply by taking as the estimate of the actual transformation the transformation corresponding to one of the coarse scale registrations. To determine which of the coarse scale registration transformations to use as the estimate for the optimal transformation, one should monitor the registration process and determine when the registration transformations begin to converge to the optimal value. In [13], we showed that in the case of rigid registration, using one of the coarse scale registration transformations to estimate the actual transformation is sufficient for a high degree of accuracy. Experimentally, we have found in [13] and in the simulations studied in this paper, that the registration transformations computed for small values of k (e.g. 2,3,4) are indeed very close to the optimal transformations. Finally, we note that the practitioner should proceed with caution in using an average or a weighted average to compute the final registration parameters; for example, the average can fail to be one-to-one, even if all of the original transformations are one-to-one.

2) Algorithm II: Multi-node registration: In our multi-node multiscale registration algorithm, Algorithm II, we register the k^{th} scale of image B with the k^{th} scale of image A , for $k = 1, 2, \dots, m$. We refer to this algorithm as a *multi-node multiscale registration algorithm* because in each of the m registrations prescribed by the algorithm, we consider both the scales of the fixed image A and the scales of the moving image B . As before, we expect that registration of $C_k(B)$ with $C_k(A)$ should be more accurate than ordinary registration for the first few values of k . As k increases, however, we expect that eventually the scales $C_k(A)$ and $C_k(B)$ will become too noisy to register successfully. Since this algorithm considers scales of both images, we expect that it will be particularly successful when both images are noisy.

Upon registering $C_k(B)$ with $C_k(A)$ for $k = 1, 2, \dots, m$, we estimate the actual transformation that maps the images to one another either by computing a weighted average of the transformations produced by these m registrations, or simply by using one of the coarse scale registrations $C_k(B)$ with $C_k(A)$ (for some small k), as in Algorithm I.

3) Algorithm III: Iterated multiscale registration: The accuracy and speed of convergence of Algorithms I and II can be improved by implementing an iterated multiscale registration algorithm, as follows. To iterate Algorithm I, we first register the first coarse scale $C_1(B)$ of the moving image with the fixed image A . The output of this registration

process is a set of deformation parameters that represent the deformation transformation produced by the registration algorithm. We then register the second scale $C_2(B)$ of the moving image with the fixed image A , using the output deformation parameters from the first registration as the starting parameters for the second registration. We repeat this procedure until the last scale (or desired stopping scale) is reached, or until the similarity measure between the images does not improve at the next scale registration. That is, at each stage, we use the output deformation parameters from the previous registration as the initial parameters for the current registration. The iterated form of Algorithm II is analogous, but at each stage we register the k^{th} scale $C_k(B)$ of the moving image with the k^{th} component $C_k(A)$ of the fixed image. We expect the single-node iterated multiscale registration algorithm to be particularly effective when only one of the images to be registered contains significant levels of noise, and we expect the multi-node iterated multiscale registration algorithm to be particularly effective when both of the images contain noise.

V. RESULTS AND DISCUSSION

A. Ordinary registration of a deformed image

Consider the mid-sagittal brain slice I and the deformed image S , shown in Figure 1. The image I is a synthesized MR-T1 mid-sagittal brain slice (217×180 pixels, 1×1 mm spacing) taken from the Insight Segmentation and Registration Toolkit (ITK) data repository [8]. The deformed image S is obtained by applying a known B-spline deformation to the original image I . To emphasize the effect of the deformation, we also illustrate in Figure 1 the difference between the deformed and original images. Since the deformation transformation that maps the deformed image S to the original image I and corresponding deformation field are known, we can effectively evaluate the accuracy of various deformable registration methods by comparing the output deformation fields with the known deformation field. Unless noted otherwise, for all registration simulations presented in this paper, we use a B-spline deformable registration technique with a mean squares image metric and a conjugate gradient descent algorithm. However, the multiscale registration algorithms developed in this paper are independent of the registration technique used to register the images.

Using a B-spline deformable registration model, the image S is successfully registered with the image I . In Figure 2, we compare the result of

the registration process, namely the image obtained upon applying the optimal deformable transformation determined by the algorithm to the deformed image, with the original image I , and we also illustrate the difference between the original and resulting images. Ideally, the original and resulting images should be identical and the difference image should be zero (all black). Indeed, the images in Figure 2 demonstrate that the deformable registration algorithm recovers the deformation transformation.

In Figure 3, we display the exact deformation field corresponding to the deformation transformation between the images I and S (on the left) and the deformation field determined by the deformable registration algorithm, and note that visually the two deformation fields are almost identical. The deformation field is a two-dimensional vector field that represents graphically the magnitude of the deformation at each point in the image. To quantitatively compare the exact and computed deformation fields, and for future reference, we compute the mean square difference (MSD) between the fields. For the deformation fields shown in Figure 2, the MSD is $7.2 \cdot 10^{-2}$. The minimum possible MSD between two fields is 0 (corresponding to no difference between the images), and the maximum possible MSD is 1 (corresponding to a large difference).

B. Ordinary registration of a noisy deformed image

In this section, we present the registration results obtained upon registering the noisy deformed images S_δ with the original image I for increasing noise variances δ . To obtain the noisy deformed images S_δ , we first deform the original image I and then add speckle noise of variance δ to the resulting deformed image. In Figure 5, we illustrate the deformation fields produced by the B-spline deformable registration algorithm upon registering the noisy deformed images S_δ with the original image I . Recall that the actual deformation is shown in Figure 3.

A visual comparison of the deformation fields presented in Figure 5 with the exact deformation field in Figure 3 indicates that the deformation registration technique fails to produce physically meaningful results for noise variance δ greater than 0.2. To quantitatively compare the deformation fields determined by the deformable registration algorithm in Figure 5 with the exact deformation field in Figure 3, we compute the mean square differences (MSD) between the computed and exact deformation fields for each speckle noise variance δ . In Table I, we present the mean square values

for each noise variance δ . For reference, we also include in the first line of Table I the MSD between the deformation field produced by the registration algorithm when the deformed image contains no noise and the exact deformation field (from Section V-A).

The results presented in Figure 5 and Table I indicate that ordinary deformable registration techniques fail to produce an accurate registration result when one of the images to be registered contains significant levels of noise. As expected, the level of failure increases as the speckle noise variance δ increases. For variances greater than or equal to 0.2 the algorithm fails to register the images.

C. Denoising results

Since ordinary deformable registration of the noisy images fails, we study next the effect of denoising on the noisy image registration problem. We apply a mean (3×3), median (3×3), and Gaussian convolution filter (kernel size 9×9 and variance 5), as discussed in Section III-B, to the noisy image $S_{0.6}$ and register the denoised images with the original image I . In Figure 6, we illustrate the denoised images and the deformation fields produced by registering the denoised images with the original image.

To quantitatively compare the computed deformation fields with the exact deformation, we compute the mean square differences between the deformation fields in Figure 6 and the exact deformation field in Figure 3. The mean square differences are $3.0 \cdot 10^{-1}$, $3.2 \cdot 10^{-1}$, and $4.6 \cdot 10^{-1}$, for mean, median, and Gaussian denoising techniques, respectively. These mean square differences, combined with a visual comparison of the deformation fields in Figure 6 with the exact deformation field in Figure 3, demonstrate that the application of classical denoising techniques prior to registration does not enable successful deformable registration of the noisy image $S_{0.6}$ with the original image I .

D. Multiscale registration of a noisy deformed image

Next, we demonstrate with several image registration experiments that the multiscale techniques presented in Section IV-B accurately and efficiently register the noisy deformed images considered in Sections V-B and V-C. Consider again the original image I and the noisy deformed image $S_{0.6}$, and recall that the deformation that maps the deformed (moving) image to the original (fixed) image is given by the deformation field in Figure 3.

TABLE I

THE MSDS BETWEEN THE COMPUTED AND EXACT DEFORMATION FIELDS FOR INCREASING NOISE VARIANCES δ FOR THE DEFORMATION FIELDS SHOWN IN FIGURE 5.

δ	0	0.1	0.2	0.3	0.4
MSD	$7.2 \cdot 10^{-3}$	$1.5 \cdot 10^{-2}$	$1.9 \cdot 10^{-2}$	$3.3 \cdot 10^{-1}$	$3.6 \cdot 10^{-1}$
δ	0.5	0.6	0.7	0.8	0.9
MSD	$4.3 \cdot 10^{-1}$	$4.4 \cdot 10^{-1}$	$4.5 \cdot 10^{-1}$	$4.8 \cdot 10^{-1}$	$5.7 \cdot 10^{-1}$

1) *Single-node multiscale registration:* In this section, we use Algorithm I to register the hierarchical scales of the noisy image $S_{0.6}$ (see Figure 7) with the original image I . The resulting deformation fields are illustrated in Figure 8.

In Table II, we calculate the MSDs between the deformation fields obtained upon registering the hierarchical scales of the noisy deformed image with the original image and the exact deformation field.

The deformation fields shown in Figure 8 and the mean square differences in Table II demonstrate that the multiscale registration method is successful for registration of the noisy deformed image $S_{0.6}$ with the original image I . Visually comparing the deformation fields in Figure 8 with the exact deformation in Figure 3, we see that the deformation fields obtained upon registering the coarse scales of the noisy deformed image with the original image are a good approximation for the actual deformation field. The MSDs indicate that the first few coarse scales produce the most accurate registration, as expected. As more detail, and hence noise, is added in increasing scales, the registration becomes less accurate for higher scales. However, it is clear both from the deformation fields in Figure 8 and the MSDs in Table II that all scale registrations of the multiscale registration technique produce very accurate results. Since both ordinary deformable registration of the noisy deformed image and the application of classical denoising algorithms prior to registration failed to produce accurate results for noise variances δ greater than 0.2, and we are able to accurately register the deformed noisy image with variance 0.6, we conclude that the multiscale single-node registration technique is a significant improvement over ordinary registration.

2) *Iterated multiscale registration:* In this section, we register the noisy deformed image $S_{0.6}$ with the original image I using (the single-node version of) Algorithm III. That is, we first decompose the noisy deformed image $S_{0.6}$ into 8 hierarchical scales using the multiscale hierarchical decomposition presented in Section IV-A. We then register

the first scale of $S_{0.6}$ with I , and use the resulting transform parameters to register the second scale of $S_{0.6}$ with I . This process is then iterated, at the last stage using the transformation parameters obtained by registering the seventh scale of $S_{0.6}$ with I as the starting parameters for registering the eighth scale of $S_{0.6}$ with I . In Table III, we compute the MSDs between the computed deformation fields using the iterated multiscale registration algorithm with the exact deformation field, and in Figure 9, we illustrate the deformation field produced by the final iteration.

The results presented in Table III and Figure 9 demonstrate that the iterated multi-scale registration algorithm is a significant improvement over ordinary deformable registration techniques. Additionally, the results of the iterated multiscale algorithm are more accurate than those obtained via Algorithm I (single-node multiscale registration), as indicated by the MSDs in Table III. However, the main improvement of the iterated algorithm over the single-node multiscale algorithm is the improvement in computation time. Working on a Dell Dimension 8400 Intel Pentium 4 CPU (3.40 GHz, 2.00 GB of RAM), registering a single scale of the noisy deformed image $S_{0.6}$ with the original image I takes an average of 84.5 seconds per scale. The first scale registration requires approximately 30 seconds, and the final scale registration requires approximately 130 seconds. The time required for the registration process increases with each scale, as each scale contains more detail and noise than the previous scale. Thus the total time required to register each scale of the noisy deformed image $S_{0.6}$ with the original image I is approximately 676 seconds. With the iterated multiscale registration method, the initial registration of the first scale of the noisy deformed image with the first scale of the original image takes approximately 29.5 seconds, and each successive iteration takes approximately 8 seconds (since the registration algorithm starts from an initial point which is very close to the optimal transformation). Thus the total time required to complete all 8 iterations of the iterated

TABLE II

THE MSDS BETWEEN THE COMPUTED AND EXACT DEFORMATION FIELDS OBTAINED UPON REGISTERING THE HIERARCHICAL SCALES OF THE NOISY IMAGE WITH THE ORIGINAL IMAGE (USING ALGORITHM I).

Scale	1	2	3	4
MSD	$1.7 \cdot 10^{-2}$	$1.1 \cdot 10^{-2}$	$9.5 \cdot 10^{-3}$	$1.0 \cdot 10^{-2}$
Scale	5	6	7	8
MSD	$1.0 \cdot 10^{-2}$	$1.1 \cdot 10^{-2}$	$9.4 \cdot 10^{-2}$	$1.0 \cdot 10^{-1}$

TABLE III

THE MSDS BETWEEN THE COMPUTED AND EXACT DEFORMATION FIELDS OBTAINED UPON REGISTERING THE NOISY DEFORMED IMAGE $S_{0.6}$ WITH THE ORIGINAL IMAGE I USING THE SINGLE-NODE ITERATED MULTISCALE REGISTRATION METHOD (ALGORITHM III).

Iteration	1	2	3	4
MSD	$1.6 \cdot 10^{-2}$	$1.1 \cdot 10^{-2}$	$6.9 \cdot 10^{-3}$	$6.4 \cdot 10^{-3}$
Iteration	5	6	7	8
MSD	$6.3 \cdot 10^{-3}$	$6.2 \cdot 10^{-3}$	$6.2 \cdot 10^{-3}$	$6.2 \cdot 10^{-3}$

multiscale registration algorithm is approximately 85.5 seconds. Hence we conclude that the iterated multiscale method is significantly more efficient than the single-node multiscale method. Moreover, the mean square differences in Table III indicate that accurate results can be achieved by iterating the iterated multiscale registration algorithm until the fourth or fifth iteration, which further reduces the computation time.

3) *Increasing the noise variance:* Finally, we demonstrate that the iterated multiscale registration algorithm produces accurate results for noise variances δ significantly greater than those at which ordinary deformable registration fails. In Figure 10, we illustrate the noisy deformed images S_δ for very large values of the noise variance δ , and in Figure 11, we illustrate the deformation fields computed using the single-node iterated multiscale registration algorithm (Algorithm III) to register the noisy deformed images S_δ with the original image I for each δ illustrated in Figure 10. In Table IV, we illustrate the MSDs between the computed and exact deformation fields. These results demonstrate that the iterated multiscale registration algorithm accurately registers the noisy deformed image with the original image for noise variances that are significantly greater than those at which ordinary registration fails. Recall from Section V-B that ordinary deformable registration of a noisy deformed image with a non-noisy fixed image fails for noise variances δ greater than 0.2. In Figure 11 and Table IV, we demonstrate that the iterated multiscale registration algorithm produces accurate results for noise variances δ as large as 6. Moreover,

the iterated multiscale registration algorithm is more accurate than ordinary B-splines deformable registration when the images contain no added speckle noise. Recall that the MSD between the exact and computed deformation field in Section V-A obtained upon registering the non-noisy deformed image S with the non-noisy original image I using B-splines deformable registration is $7.2 \cdot 10^{-3}$. Using the iterated multiscale registration algorithm, the MSD between the computed and exact deformations fields is $1.8 \cdot 10^{-3}$. Thus the iterated multiscale algorithm improves the accuracy of deformable registration even when the images do not contain added noise.

E. Registration of a noisy deformed image with a noisy fixed image

In this section, we consider the case in which both images to be registered contain significant levels of noise. We add speckle noise of variance 0.6 to the original image I , and denote this noisy image $I_{0.6}$. Our goal is to register the noisy deformed image $S_{0.6}$ with the noisy fixed image $I_{0.6}$. In Figure 12, we illustrate both of the noisy images, as well as the deformation field produced upon registering the noisy deformed image $S_{0.6}$ with the noisy original image $I_{0.6}$ using ordinary B-splines deformable registration techniques.

A visual comparison of the computed deformation field in Figure 12 with the exact deformation field in Figure 3 indicates that ordinary deformable registration of the noisy images fails. The MSD in this case is $4.4 \cdot 10^{-1}$.

1) *Multi-node registration:* Since ordinary deformable registration of the noisy images fails, we

TABLE IV

THE MSDS BETWEEN THE COMPUTED AND EXACT DEFORMATION FIELDS OBTAINED UPON REGISTERING THE NOISY DEFORMED IMAGES S_δ WITH THE ORIGINAL IMAGE I FOR INCREASING NOISE VARIANCES δ USING THE SINGLE-NODE ITERATED MULTISCALE REGISTRATION METHOD (ALGORITHM III).

δ	0	0.4	0.8	1
MSD	$1.8 \cdot 10^{-3}$	$4.2 \cdot 10^{-3}$	$9.9 \cdot 10^{-3}$	$7.2 \cdot 10^{-2}$
δ	2	3	4	6
MSD	$1.4 \cdot 10^{-2}$	$1.3 \cdot 10^{-2}$	$1.7 \cdot 10^{-2}$	$6.4 \cdot 10^{-2}$

register the images using our multiscale multi-node algorithm (Algorithm II). That is, we apply the multiscale decomposition to both of the images $I_{0.6}$ and $S_{0.6}$ and register the k -th scale of $S_{0.6}$ with the k -th scale of $I_{0.6}$ for $k = 1, 2, \dots, 8$ (we use $m = 8$ hierarchical steps in the multiscale decomposition of each image). The deformation fields obtained upon registering the scales of $S_{0.6}$ with the scales of $I_{0.6}$ are illustrated in Figure 13. A visual comparison of the deformation fields in Figure 13 with the exact deformation field in Figure 3 indicates that the multi-node registration algorithm is successful. The corresponding MSDs are shown in Table V.

2) *Iterated multiscale registration:* In this section, we register the noisy deformed image $S_{0.6}$ with the noisy original image $I_{0.6}$ using (the multi-node version of) Algorithm III. That is, we first decompose the noisy deformed image $S_{0.6}$ and the noisy original image $I_{0.6}$ into 8 hierarchical scales each, using the multiscale hierarchical decomposition presented in Section IV-A. We then register the first scale of $S_{0.6}$ with the first scale of $I_{0.6}$, and use the resulting transform parameters to register the second scale of $S_{0.6}$ with the second scale of $I_{0.6}$. This process is then iterated, at the last stage using the transformation parameters obtained by registering the seventh scale of $S_{0.6}$ with the seventh scale of $I_{0.6}$ as the starting parameters for registering the eighth scale of $S_{0.6}$ with the eighth scale of $I_{0.6}$. In Table VI, we compute the MSDs between the computed deformation fields using the iterated multiscale registration algorithm with the exact deformation field, and in Figure 14, we illustrate the deformation field produced by the final iteration.

The results presented in Table VI and Figure 14 demonstrate that the iterated multi-scale registration algorithm is a significant improvement over ordinary deformable registration techniques. Additionally, the results of the iterated multiscale algorithm are more accurate than those obtained via Algorithm II (multi-node multiscale registration), as indicated

by the MSDs in Table VI. Additionally, the iterated algorithm is computationally more efficient than Algorithm II.

3) *Increasing the noise variance:* Finally, we demonstrate as in Section V-D.3 that the iterated multiscale registration algorithm produces accurate results when both of the images contain speckle noise of variance significantly greater than the level at which ordinary deformable registration fails. In Figure 16, we illustrate the deformation fields computed using the iterated multiscale registration algorithm to register the noisy deformed image S_δ with the noisy original image I_δ for increasing noise variances δ . In Table VII, we illustrate the mean square differences between the computed and exact deformation fields. These results demonstrate that the iterated multiscale registration algorithm accurately registers the noisy deformed image with the noisy original image for noise variances significantly greater than those at which ordinary techniques fail; recall that ordinary deformable registration failed when *only one* of the images to be registered contain noise of variance 0.2. In Figure 15, we illustrate the noisy original and deformed images I_2 and S_2 . These images contain speckle noise with variance $\delta = 2$. As demonstrated by the deformation field in Figure 16, the multi-node iterated multiscale registration algorithm (Algorithm III) accurately registers these very noisy images.

F. Finite element registration results

In this section, we demonstrate that the multiscale registration technique works accurately in conjunction with other deformable registration methods. In this section, we demonstrate that the multiscale registration technique works for a more general class of examples (i.e. our technique is not dependent on the type of deformation model assumed or on the type of deformable registration algorithm used). In particular, we shall demonstrate here that the multiscale technique works in conjunction with finite element deformable registration. Finite element deformable registration models [7] use biomechanical models

TABLE V

THE MSDS BETWEEN THE COMPUTED AND EXACT DEFORMATION FIELD OBTAINED UPON REGISTERING THE HIERARCHICAL SCALES OF THE NOISY DEFORMED IMAGE $S_{0.6}$ WITH THE HIERARCHICAL SCALES OF THE NOISY ORIGINAL IMAGE $I_{0.6}$ (USING ALGORITHM II).

Scale	1	2	3	4
MSD	$1.2 \cdot 10^{-2}$	$1.2 \cdot 10^{-2}$	$1.0 \cdot 10^{-2}$	$1.1 \cdot 10^{-2}$

Scale	5	6	7	8
MSD	$1.2 \cdot 10^{-2}$	$1.7 \cdot 10^{-2}$	$1.8 \cdot 10^{-2}$	$1.9 \cdot 10^{-2}$

TABLE VI

THE MSDS BETWEEN THE COMPUTED AND EXACT DEFORMATION FIELD OBTAINED UPON REGISTERING THE NOISY DEFORMED IMAGE $S_{0.6}$ WITH THE NOISY ORIGINAL IMAGE $I_{0.6}$ USING THE MULTI-NODE ITERATED MULTISCALE REGISTRATION METHOD (ALGORITHM III).

Iteration	1	2	3	4
MSD	$1.3 \cdot 10^{-2}$	$1.2 \cdot 10^{-2}$	$1.0 \cdot 10^{-2}$	$1.0 \cdot 10^{-2}$

Iteration	5	6	7	8
MSD	$1.0 \cdot 10^{-2}$	$9.5 \cdot 10^{-3}$	$9.0 \cdot 10^{-3}$	$8.0 \cdot 10^{-3}$

TABLE VII

THE MSDS BETWEEN THE COMPUTED AND EXACT DEFORMATION FIELDS OBTAINED UPON REGISTERING THE NOISY DEFORMED IMAGES S_δ WITH THE NOISY IMAGE I_δ FOR INCREASING NOISE VARIANCES δ USING THE MULTI-NODE ITERATED MULTISCALE REGISTRATION METHOD (ALGORITHM III).

δ	0	0.4	0.8	1	1.5	2
MSD	$1.8 \cdot 10^{-3}$	$9.1 \cdot 10^{-3}$	$3.7 \cdot 10^{-2}$	$4.4 \cdot 10^{-2}$	$2.0 \cdot 10^{-2}$	$9.1 \cdot 10^{-2}$

for deformable brain registration in which the brain is modeled as an elastic body with an interstitial fluid. As our goal here is primarily to demonstrate that our multiscale registration algorithm can be used in conjunction with *any* deformable registration model, we do not present here the details of finite element registration and instead refer the reader to [7].

We again consider registration of the noisy deformed images S_δ with the original image I for increasing values of the noise variance δ . Using the iterated single-node multiscale registration algorithm (Algorithm III) *in conjunction with a finite element deformable registration technique*, we register the noisy deformed images S_δ with I . In Table VIII, we present the MSD values between the computed and exact known deformation fields. Comparing these MSD values with those presented, for example, in Table IV, we conclude that the multiscale registration algorithm works as accurately with this different deformable registration model as it does with a B-spline deformable registration model.

G. Multi-modality registration results

In this section, we demonstrate that the multiscale registration technique accurately registers noisy images of different modalities. To study the multi-modality problem, we use one brain proton density slice image and one brain T1 slice image, as illustrated in Figure 17. Both images are taken from the ITK data repository [8]. To study the problem of multi-modality image registration in the presence of noise, we first deform the T1 slice B using the known B-spline deformation illustrated in Figure 3, and then add speckle noise of increasing variance δ to the resulting deformed image. We denote by P the original brain proton density image and by T_δ the deformed T1 image with noise of variance δ . We first register the noisy deformed images T_δ with the original image P for increasing values of δ using ordinary B-spline deformable registration techniques *with a Viola-Wells mutual information image metric*. In the case of registering images of different modalities, metrics based on direct comparison of gray levels are generally not applicable, and it has been shown extensively that metrics based on mutual information are well-suited for multi-modality registration. In Table IX, we present the

TABLE VIII

THE MSDS BETWEEN THE COMPUTED AND EXACT DEFORMATION FIELDS OBTAINED UPON REGISTERING THE NOISY DEFORMED IMAGES S_δ WITH THE ORIGINAL IMAGE I FOR INCREASING NOISE VARIANCES δ USING THE SINGLE-NODE ITERATED MULTISCALE REGISTRATION METHOD (ALGORITHM III) IN CONJUNCTION WITH A FINITE ELEMENT REGISTRATION TECHNIQUE.

δ	0	0.2	0.4	0.6	0.8	1.0
MSD	$1.9 \cdot 10^{-3}$	$2.7 \cdot 10^{-3}$	$4.1 \cdot 10^{-3}$	$5.9 \cdot 10^{-3}$	$1.1 \cdot 10^{-2}$	$7.4 \cdot 10^{-2}$

mean square values between the deformation field computed by the B-splines deformable registration algorithm and the exact known deformation field. The MSD values presented in Table IX demonstrate that ordinary B-splines registration fails to accurately register the images when one of the images to be registered contains significant levels of noise. This failure can be seen by comparing the MSD values presented here to those presented in Table I and looking at the corresponding deformation fields in Figure 5, where we demonstrate the failure of ordinary B-spline registration for the registration of the noisy images S_δ with I .

Next, we demonstrate that our multiscale registration algorithm enables successful registration of the multi-modality noisy images. In Table X, we present the MSD values between the computed and exact deformation field obtained upon registering the noisy deformed images T_δ with P using the single-node iterated multiscale registration method (Algorithm III) in conjunction with a B-splines deformable registration algorithm with a Viola-Wells mutual information image metric. The MSD values in Table X illustrate the accuracy of the multiscale registration algorithm for noisy images of different modalities.

Finally, we note that the multi-node iterated multiscale registration technique accurately registers the proton density and T1 images when both of the images contain noise. We do not include the results here in order to keep the paper concise.

H. Extension to three-dimensional images

Finally, we demonstrate that the multiscale registration technique accurately registers *three-dimensional* images. To study the three-dimensional (3D) registration problem, we use four-dimensional computed tomography (4D CT) images acquired with a GE Discovery-ST Scanner (GE Medical Systems, Milwaukee, WI) at the Stanford University Medical Center. Four-dimensional computed tomography techniques allows one to acquire image data at specified phases over several respiratory cycles, and then combines the data into

three-dimensional phase-binned images. We obtained eight phase bins (i.e. eight 3D CT images) corresponding to eight different breathing phases of the respiratory cycle. Each phase consists of 80 two-dimensional images, or *slices*, which are combined to obtain the 3D images. The slice thickness for each phase is 2.5-mm, and the eight breathing phases recorded contain approximately 400 MB of data in DICOM image format.

In Figure 18, we illustrate two corresponding sample slices (slice 25) from the first and eighth phases of the 4D CT data set.

To register the 3D CT images with one another (i.e. to register each phase of the respiratory cycle with, for example, the inhale phase), we first extend the hierarchical multiscale decomposition of [19] to 3D images. Although the multiscale decomposition presented in [19] was done in two dimensions only, the hierarchical multiscale expansion in Eq. (6) is independent of the image dimensionality. To implement the iterated multiscale decomposition in 3 dimensions, we cover the image domain Ω with a grid ($x_i := ih, y_j := jh, z_k := kh$), and let D_+ , D_- , and D_0 denote the forward, backward, and centered divided differences, respectively. Then the 3D extension of the iterated multiscale decomposition given by Eq. (12) in Section IV-A.1 is:

TABLE IX

THE MSDS BETWEEN THE COMPUTED AND EXACT DEFORMATION FIELDS OBTAINED UPON REGISTERING THE NOISY DEFORMED IMAGES T_δ WITH THE ORIGINAL IMAGE P FOR INCREASING NOISE VARIANCES δ USING ORDINARY B-SPLINE DEFORMABLE REGISTRATION.

δ	0	0.2	0.4	0.6	0.8	1.0
MSD	$4.9 \cdot 10^{-3}$	$3.2 \cdot 10^{-2}$	$8.7 \cdot 10^{-2}$	$2.1 \cdot 10^{-1}$	$3.8 \cdot 10^{-1}$	$6.4 \cdot 10^{-1}$

TABLE X

THE MSDS BETWEEN THE COMPUTED AND EXACT DEFORMATION FIELDS OBTAINED UPON REGISTERING THE NOISY DEFORMED T1 IMAGES T_δ WITH THE ORIGINAL PROTON DENSITY IMAGE P FOR INCREASING NOISE VARIANCES δ USING THE SINGLE-NODE ITERATED MULTISCALE REGISTRATION METHOD (ALGORITHM III).

δ	0	0.2	0.4	0.6	0.8	1.0
MSD	$1.6 \cdot 10^{-3}$	$2.5 \cdot 10^{-3}$	$4.3 \cdot 10^{-3}$	$6.7 \cdot 10^{-3}$	$1.0 \cdot 10^{-2}$	$8.1 \cdot 10^{-2}$

$$\begin{aligned}
u_{i,j,k}^{n+1} &= f_{i,k,j} \\
&+ \frac{1}{2h^2} \left[\frac{u_{i+1,j,k}^n - u_{i,j,k}^{n+1}}{\sqrt{\epsilon^2 + (D_{+x}u_{i,j,k}^n)^2 + (D_{0y}u_{i,j,k}^n)^2 + (D_{0z}u_{i,j,k}^n)^2}} \right] \\
&- \frac{1}{2h^2} \left[\frac{u_{i,j,k}^{n+1} - u_{i-1,j,k}^n}{\sqrt{\epsilon^2 + (D_{-x}u_{i,j,k}^n)^2 + (D_{0y}u_{i-1,j,k}^n)^2 + (D_{0z}u_{i-1,j,k}^n)^2}} \right] \\
&+ \frac{1}{2h^2} \left[\frac{u_{i,j+1,k}^n - u_{i,j,k}^{n+1}}{\sqrt{\epsilon^2 + (D_{0x}u_{i,j,k}^n)^2 + (D_{+y}u_{i,j,k}^n)^2 + (D_{0z}u_{i,j,k}^n)^2}} \right] \\
&- \frac{1}{2h^2} \left[\frac{u_{i,j,k}^{n+1} - u_{i,j-1,k}^n}{\sqrt{\epsilon^2 + (D_{0x}u_{i,j-1,k}^n)^2 + (D_{-y}u_{i,j,k}^n)^2 + (D_{0z}u_{i,j-1,k}^n)^2}} \right] \\
&+ \frac{1}{2h^2} \left[\frac{u_{i,j,k+1}^n - u_{i,j,k}^{n+1}}{\sqrt{\epsilon^2 + (D_{0x}u_{i,j,k}^n)^2 + (D_{0y}u_{i,j,k}^n)^2 + (D_{+z}u_{i,j,k}^n)^2}} \right] \\
&- \frac{1}{2h^2} \left[\frac{u_{i,j,k}^{n+1} - u_{i,j,k-1}^n}{\sqrt{\epsilon^2 + (D_{0x}u_{i,j,k-1}^n)^2 + (D_{0y}u_{i,j,k-1}^n)^2 + (D_{-z}u_{i,j,k}^n)^2}} \right].
\end{aligned} \tag{15}$$

Upon decomposing each 3D image, we use the iterated multi-node multiscale registration algorithm (Algorithm III) to register each scale of the moving 3D image with each scale of the fixed 3D image. In Figure 19, we compare the voxel-wise intensity difference between the two sample slices shown in Figure 18 before and after iterated multiscale deformable registration. In these images, black represents exact intensity agreement and brighter regions indicate intensity disagreement. A comparison of the intensity differences before and after registration demonstrates that the iterated multiscale registration method indeed recovers the difference between the two images. Similar results are obtained with all other slices and phases. To quantitatively evaluate the accuracy of the 3D multiscale method, we compare the MSD values between the 3D images before registration and after ordinary and multiscale deformable registration. The MSD values are $9.1 \cdot 10^{-1}$ before registration, $8.2 \cdot 10^{-4}$ after ordinary deformable registration, and $7.6 \cdot 10^{-4}$ after multiscale registration. Working on a Dell Dimension 8400 Intel Pentium 4 CPU (3.40 GHz, 2.00 GB of RAM), the total required computation time for both the 3D multiscale decomposition and iterated multiscale multi-node registration algorithm is on the order of approximately 15-30 minutes, depending on the data set. The quality of the results of the 3D multiscale registration is similar to that of the results of the 2D multiscale registration, both for noisy and non-noisy images. The 3D multiscale registration technique enables accurate registration of noisy 3D images, exactly as in the 2D case. We do not show such examples here in order to keep the paper concise.

VI. CONCLUSIONS

While there are many existing deformable registration techniques, common approaches are shown to fail when one or more of the images to be registered contains even moderate levels of noise. Further, for high levels of noise, image registration fails even when classical denoising algorithms are applied to the images before registration. We have presented deformable image registration techniques based on the hierarchical multiscale image decomposition of [19]. The multiscale decomposition of an image results in a hierarchical representation that separates the coarse and fine scales of the image. We presented three separate multiscale registration algorithms based on this decomposition. In the first, we follow a single-node multiscale registration strategy in which we register the scales of the moving image with the fixed image, and use a

weighted average to estimate the actual deformation between the images. In the second, we use a multi-node multiscale registration method in which we register the scales of the moving image with the scales of the fixed image, and use a weighted average to estimate the actual deformation between the images. Finally, in the third algorithm, we follow an iterated multiscale registration strategy. Using images in which the precise deformation between the fixed and moving images is known, we have shown that the multiscale registration algorithms are indeed accurate for levels of noise much higher than the noise levels at which ordinary deformable registration techniques fail. Moreover, the iterated multiscale registration technique significantly reduces the computation time necessary to obtain accurate registration of noisy images. We have also demonstrated that the multiscale techniques improve the accuracy of ordinary B-splines deformable registration even when the images to be registered do not contain additional noise. Additionally, we have extended the hierarchical multiscale decomposition of [19] to three-dimensional images, and have demonstrated the accuracy and efficiency of our multiscale registration techniques for the registration of three-dimensional images. Although we have presented our algorithm in a way that is, in principle, independent of the specific multiscale decomposition used for the expansion of the images to be registered, we have found that the hierarchical (BV, L^2) multiscale decomposition of [19] contains unique features that are not necessarily evident in other decomposition techniques. For example, information about small geometrical details is contained in both the coarse and fine scales of the image decomposition (see, for example, Figure 7). For further details, we refer to [19]. Thus we conclude that the hierarchical (BV, L^2) decomposition is particularly well-suited for image registration problems. Finally, we would like to emphasize that using the multiscale decomposition is independent of the registration method used and of the noise model. The multiscale decomposition can be used in conjunction with any registration method and can be applied to registration of images containing any type of noise, without any assumption about the particular type of noise contained in the images.

ACKNOWLEDGMENT

The work of D. Levy was supported in part by the National Science Foundation under Career Grant No. DMS-0133511. The work of L. Xing was supported in part by the Department of Defense under Grant No. PC040282 and the National Cancer

Institute under Grant No. 1R01 CA104205, and the American Cancer Society Grant No. RSG-01-022-01-CCE.

REFERENCES

- [1] A. ACHIM, A. BEZERIANOS, AND P. TSAKALIDES, *Novel Bayesian Multiscale Method for Speckle Removal in Medical Ultrasound Images*, IEEE Transactions on Medical Imaging, **20**, no. 8, pp. 772–783, 2001.
- [2] R. BAJCSY, ET AL., *A computerized system for the elastic matching of deformed radiographic images to idealized atlas images*, Journal of Computer Assisted Tomography, **7**, no. 4, pp. 618–625, 1983.
- [3] F. L. BOOKSTEIN, *Principal Warps: Thin-Plate Splines and the Decomposition of Deformations*, IEEE Transactions on Pattern Analysis and Machine Intelligence, **11**, no. 6, pp. 567–585, 1989.
- [4] L. BROWN, *A Survey of Image Registration Techniques*, ACM Computing Surveys, **24**, no. 4, pp. 325–376, 1992.
- [5] G. E. CHRISTENSEN, R. D. RABBITT, AND M. I. MILLER, *Deformable templates using large deformation kinematics*, IEEE Transactions on Image Processing, **5**, pp. 1435–1447, 1996.
- [6] W. R. CRUM, T. HARTKENS, AND D. L. G. HILL, *Non-rigid image registration: theory and practice*, The British Journal of Radiology, **77**, pp. 140–153, 2004.
- [7] M. FERRANT, A. NAVAVI, B. MACO, P. M. BLACK, F. A. JOLESZ, R. KIKINIS, ET AL., *Serial registration of intraoperative MR images of the brain*, Medical Image Analysis, **6**, pp. 337–359, 2002.
- [8] *Insight Segmentation and Registration Toolkit (ITK)*, <http://www.itk.org>.
- [9] S. LEE, G. WOLBERG, K.-Y. CHWA, AND S. Y. SHIN, *Image metamorphosis with scattered feature constraints*, IEEE Transactions on Visualization and Computer Graphics, **2**, pp. 337–354, 1996.
- [10] S. LEE, G. WOLBERG, AND S. Y. SHIN, *Scattered data interpolation with multilevel B-splines*, IEEE Transactions on Visualization and Computer Graphics, **3**, pp. 228–244, 1997.
- [11] J. LIAN, L. XING, S. HUNJAN, ET AL., *Integrating deformable MRI/MRSI and CT image registration into the prostate IMRT treatment planning*, Medical Physics, **31**, no. 11, pp. 3087–3094, 2004.
- [12] J. MODERSITZKI, *Numerical Methods for Image Registration*, Oxford, 2004.
- [13] D. PAQUIN, D. LEVY, E. SCHREIBMANN, AND L. XING, *Multiscale image registration*, Mathematical Biosciences and Engineering, **3**, no. 2, pp. 389–418, 2006.
- [14] D. W. ROBERTS, A. HARTOV, F.E. KENNEDY, M.I. MIGA, AND K.D. PAULSEN, *Intraoperative brain shift and deformation: A quantitative analysis of cortical displacement in 28 cases*, Neurosurgery, **43**, no. 4, pp. 749–760, 1998.
- [15] L. RUDIN, S. OSHER, AND E. FATEMI, *Nonlinear total variation based noise removal algorithms*, Physica D, **60**, pp. 259–268, 1992.
- [16] D. RUECKERT, L. I. SONODA, C. HAYES, D. L. G. HILL, M. O. LEACH, AND D. J. HAWKES, *Nonrigid registration using free-form deformations: application to breast MR images*, IEEE Transactions on Medical Imaging, **18**, no. 8, pp. 712–721, 1999.
- [17] E. SCHREIBMANN, G. T.Y. CHEN, AND L. XING, *Image interpolation in 4D CT using a b-spline deformable registration model*, International Journal of Radiation Oncology, Biology, Physics, **64**, no. 5, pp. 1537–1550.
- [18] E. SCHREIBMANN, AND L. XING, *Narrow band deformable registration of prostate MRI/MRSI and CT studies*, International Journal of Radiation Oncology, Biology, Physics, **62**, no. 2, pp. 595–605, 2005.
- [19] E. TADMOR, S. NEZZAR, AND L. VESE, *A multiscale image representation using hierarchical (BV, L^2) decompositions*, Multiscale Modeling and Simulations, **2**, no. 4, pp. 554–579, 2004.
- [20] A. TEI, A. BHARATHA, M. FERRANT, P. MCL. BLACK, F.A. JOLESZ, R. KIKINIS, AND S. WARFIELD, *Tracking volumetric brain deformation during image guided neurosurgery*, VISIM: Information Retrieval and Exploration in Large Medical Image Collections, Springer Verlag, 2001.
- [21] A. R. WEEKS, *Fundamentals of Electronic Image Processing*, SPIE Optical Engineering Press and IEEE Press, 1996.

APPENDIX I
FIGURES

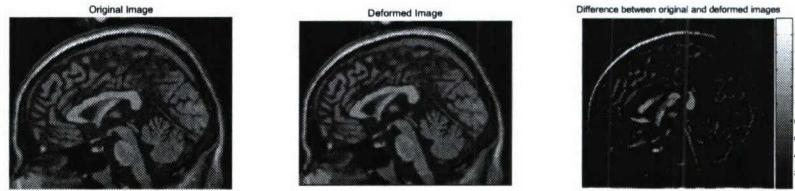


Fig. 1. The mid-sagittal brain slice I (shown on the left) and the deformed image S (shown in the center), and the difference between the original and deformed images (shown on the right).

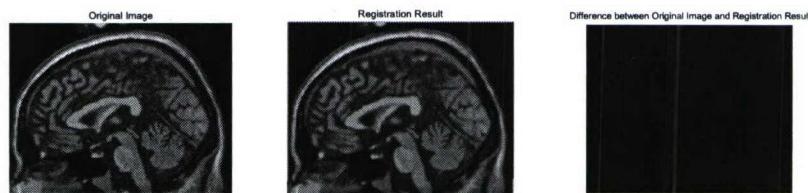


Fig. 2. The result (shown in the center) upon registering the deformed image S with the original image I (shown on the left), as well as the difference between the images (shown on the right).

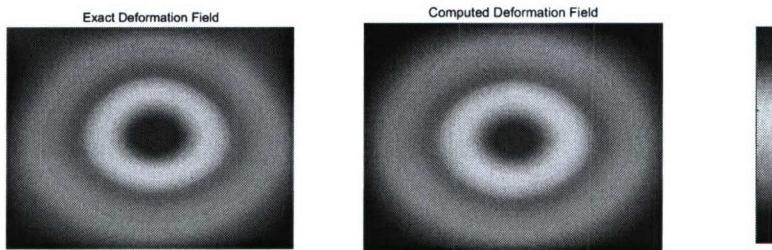


Fig. 3. The exact deformation field corresponding to the deformation transformation between I and S (shown on the left) and the deformation field produced by the registration algorithm upon registering the deformed image S with the original image I (shown on the right).

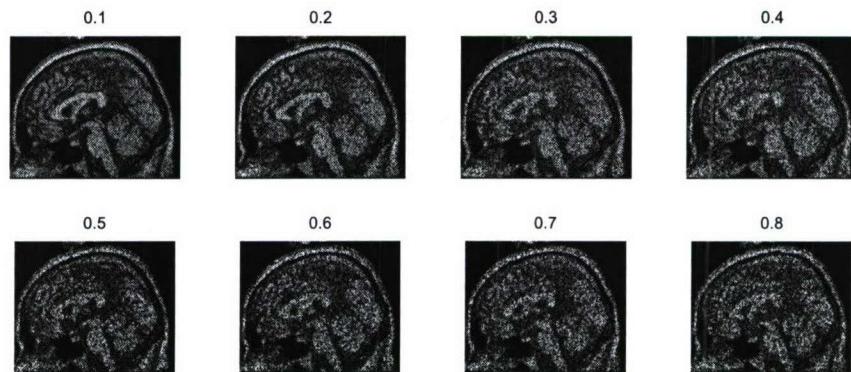


Fig. 4. The noisy images S_δ , for increasing values of δ .

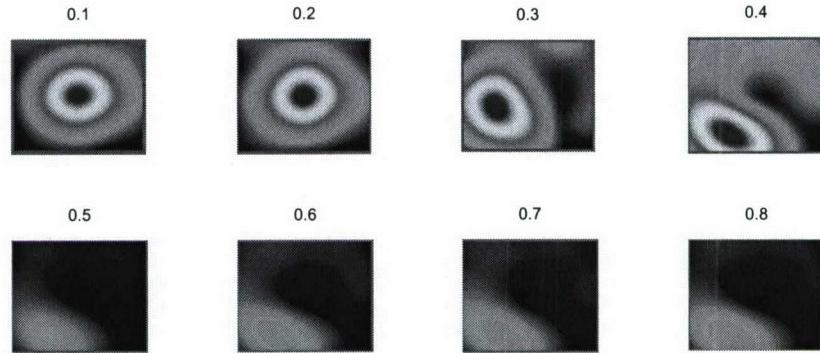


Fig. 5. The deformation fields produced by the registration algorithm upon registering the noisy deformed images S_δ with the original image I , for increasing values of δ .

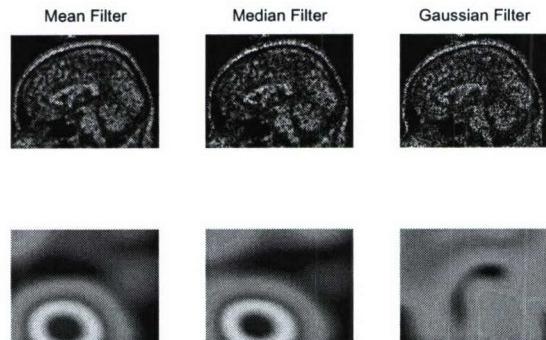


Fig. 6. The denoised images and registration results obtained upon applying mean, median, and Gaussian filters prior to registering the noisy deformed image $S_{0.6}$ with the original image I .

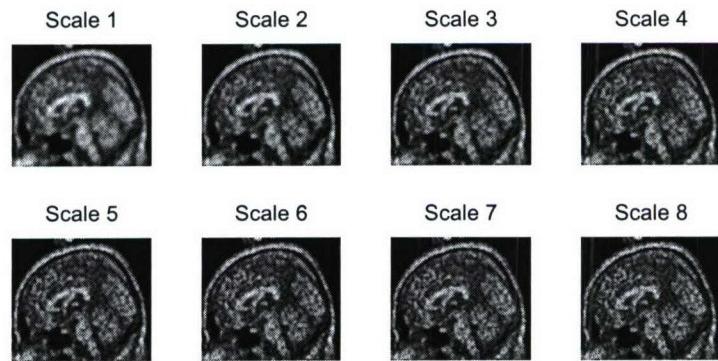


Fig. 7. The multiscale decomposition of the noisy deformed image $S_{0.6}$.

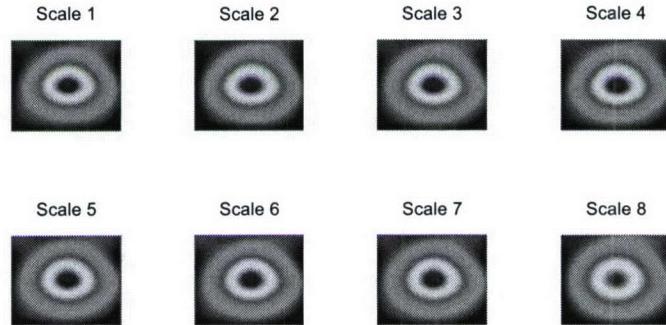


Fig. 8. The deformation fields obtained upon registering the hierarchical scales of the noisy image $S_{0.6}$ with the original image using Algorithm I.

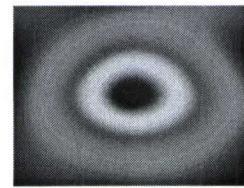


Fig. 9. The deformation field obtained upon registering the noisy deformed image $S_{0.6}$ with the original image using the single-node iterated multiscale registration method (Algorithm III).

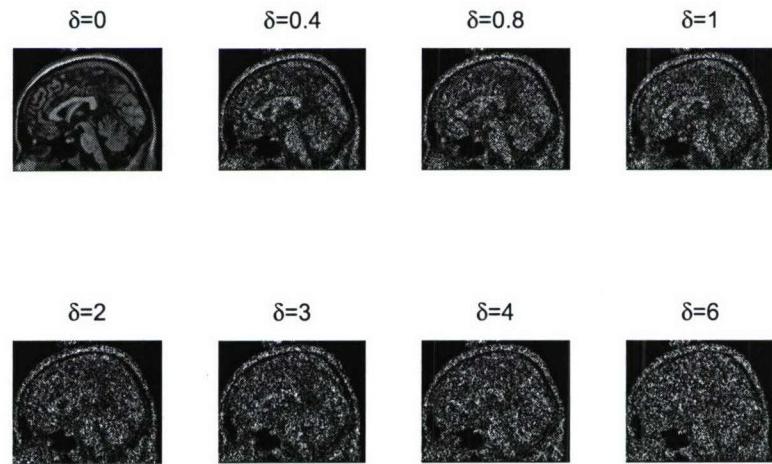


Fig. 10. The noisy deformed images S_δ for increasing noise variances δ .

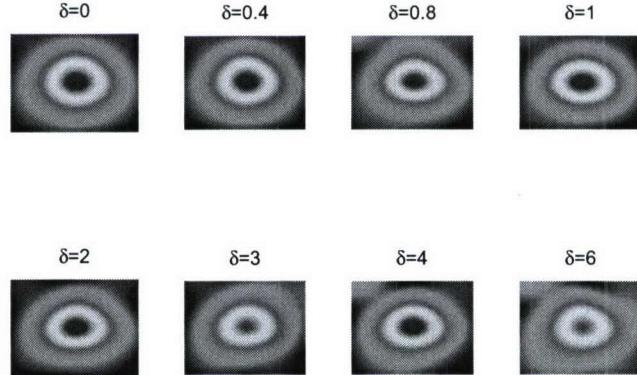


Fig. 11. The deformation fields obtained upon registering the noisy deformed image S_δ with the original image I using Algorithm III for increasing noise variances δ .

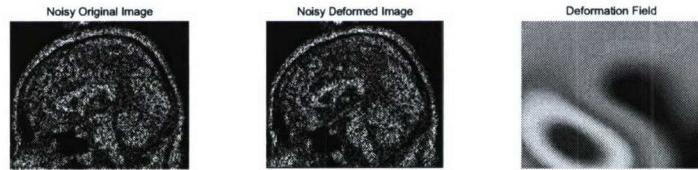


Fig. 12. The noisy mid-sagittal brain slice $I_{0.6}$ (shown on the left), the noisy deformed image $S_{0.6}$ (shown in the center), and the deformation field (shown on the right) produced upon registering $S_{0.6}$ with $I_{0.6}$ using ordinary deformable registration techniques.

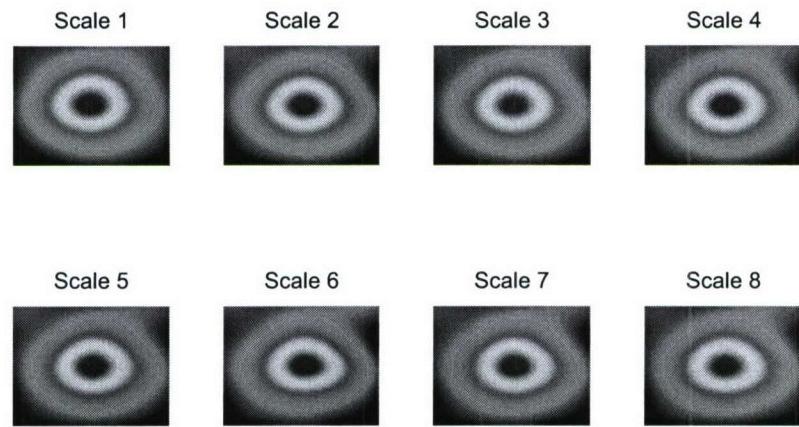


Fig. 13. The deformation fields obtained upon registering the noisy deformed image with the noisy fixed image using the multi-node registration algorithm (Algorithm II).

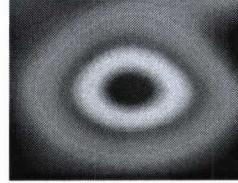


Fig. 14. The deformation field obtained upon registering the noisy deformed image $S_{0.6}$ with the noisy original image $I_{0.6}$ using the multi-node iterated multiscale registration method (Algorithm III).

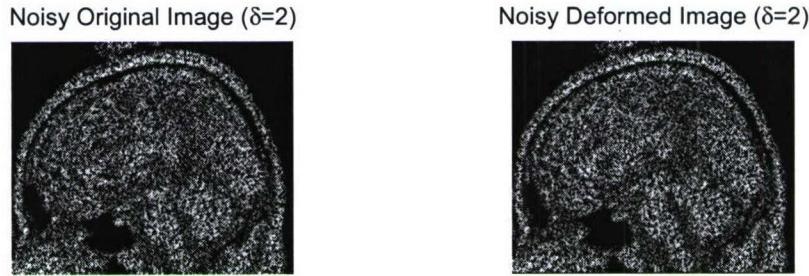


Fig. 15. The noisy original and deformed images I_2 and S_2 .

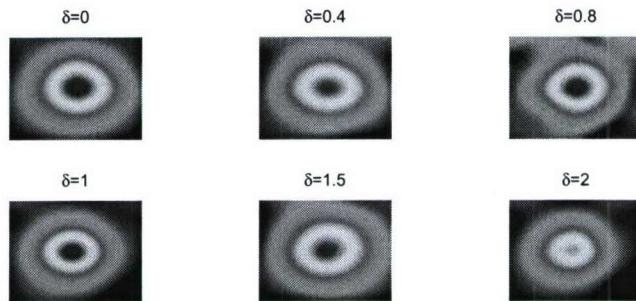


Fig. 16. The deformation fields obtained upon registering the noisy deformed image S_δ with the noisy original image I_δ using the multi-node iterated multiscale algorithm (Algorithm III) for increasing noise variances δ .



Fig. 17. The brain proton density slice (shown on the left) and T1 slice (shown on the right) images used in our multi-modality simulations.

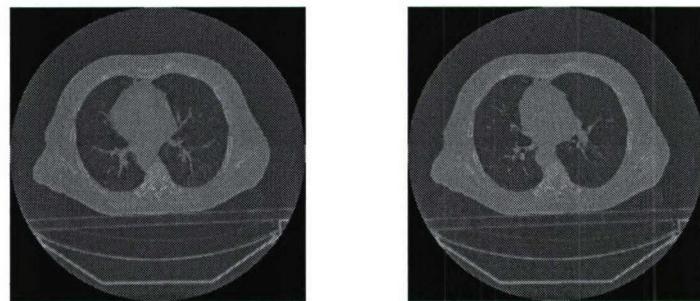


Fig. 18. Two corresponding sample slices from two breathing phases of the same patient.

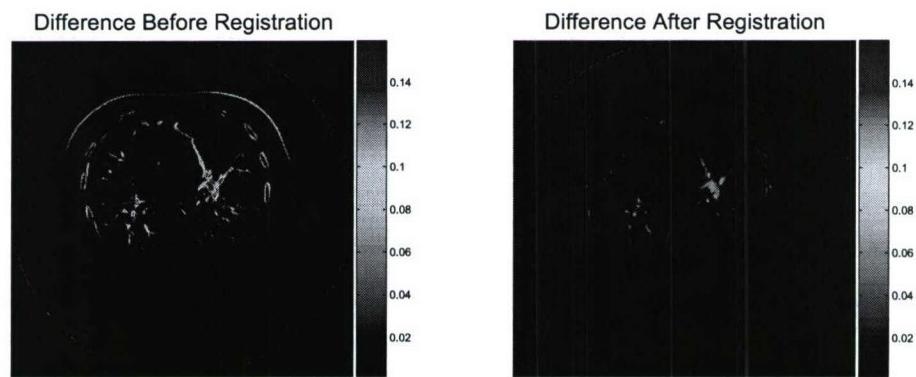


Fig. 19. The voxel-wise intensity difference between two corresponding slices before and after iterated multiscale deformable registration.

Hybrid Multiscale Landmark and Deformable Image Registration

Dana Paquin*, Doron Levy† and Lei Xing‡

July 23, 2006

Abstract

A novel hybrid image registration algorithm is presented based on combining landmark registration techniques, deformable registration techniques, and a multiscale decomposition of the images to be registered. The construction of the algorithm is motivated by the idea that the accuracy of standard deformable registration algorithms can be improved by incorporating more information, such as the location of corresponding features, about the images to be registered. The images are decomposed into coarse and fine scales, and the coarse scales are registered with one another using a landmark-based registration algorithm. The resulting transformation is then used as an initial guess to deformably register the original images with another. Using four-dimensional computed tomography (4D CT) images, the hybrid algorithm is shown to be more accurate and robust than standard deformable registration algorithms, standard landmark-based registration algorithms, combined landmark-based and deformable registration algorithms, and our previous multiscale deformable registration algorithms. Numerous image registration experiments are presented in both two and three dimensions. The hybrid algorithm is shown to be particularly robust with respect to the presence of noise in the images to be registered, as well as with respect to the location and number of landmarks used in the coarse scale landmark-based registration stage of the algorithm. The main advantages of our hybrid algorithm over traditional image registration techniques are its improved accuracy, efficiency, and robustness.

1 Introduction

Image registration is the process of determining the optimal spatial transformation that maps one image to another. Image registration is necessary, for example, when images of the same object are taken at different times, from different imaging devices, or from different perspectives. The two images to be registered with one another, called the fixed and moving images, are the input to the registration algorithm, and the output is the optimal transformation that maps the moving image to the fixed image. Ideally, the transformed moving image should be identical to the fixed image after registration. Applications of image registration include image-guided radiation therapy (IGRT), image-guided surgery, functional MRI analysis, and tumor detection, as well as many non-medical

*Department of Mathematics, Stanford University, Stanford, CA 94305-2125; paquin@math.stanford.edu

†Department of Mathematics, Stanford University, Stanford, CA 94305-2125; dlevy@math.stanford.edu

‡Department of Radiation Oncology, Stanford University, Stanford, CA 94305-5847; lei@reyes.stanford.edu

applications, such as computer vision, pattern recognition, and remotely sensed data processing (see [4] and the references therein).

Image registration models are classified into two main categories according to the transformation type: rigid and deformable. Rigid image registration models assume that the transformation that maps the moving image to the fixed image consists only of translations and rotations, while deformable models allow localized stretching of images. While rigid models are sufficient in certain circumstances, many registration problems, particularly in medical applications, are non-rigid, as most of the organs in the human body are not confined to rigid motion only. For example, respiratory motion causes non-rigid, or deformable, distortion of the lungs and other organs. As another example, image-guided neurosurgery procedures require deformable registration of pre- and intraoperative images of the brain [15], [22]. For additional applications of deformable registration, see [10], [18], [19], [20], and the references therein. Most of the current research in image registration is focused on the problem of deformable registration, and that is our focus in this paper.

Due to its importance in various applications and to the difficult nature of the registration problem, numerous image registration techniques have been proposed. See [12] and [4] for an overview of image registration techniques. Commonly used registration methods in current research are landmark-based [16], principal axes-based [1], optimal parametric [4], splines [6], elastic [2], fluid [5], and finite element [8] techniques.

In this paper, we shall be particularly interested in landmark-based and splines-based deformable registration techniques. Landmark-based registration techniques use the correspondence of a set of features, or landmarks, in the images to determine the transformation that maps the moving image to the fixed image. Although landmark-based techniques are easy to implement, the identification of corresponding features in the images to be registered is a difficult and time-consuming task, and the accuracy of such techniques is dependent on precise correspondence between landmarks [16], [20]. Splines-based deformable registration techniques are based on minimizing the distance, or metric, between the fixed and moving images. These algorithms use a grid of control points in the images to be registered to compute the optimal transformation that brings the moving image into alignment with the fixed image. Splines-based algorithms are particularly important in medical image registration, as they allow for localized deformations in the images.

Recently, we have introduced multiscale image registration techniques [13], [14] for the problem of image registration in the presence of noise. In [13], we considered the problem of *rigid* registration of noisy images, and in [14], we considered the problem of *deformable* registration of noisy images. Our multiscale registration algorithm is implemented as follows. Using the multiscale image decomposition of Tadmor, *et al.* [21], an extension of the method proposed by Rudin, *et al.* [17], we first decompose the noisy images to be registered into a series of coarse and fine scales. For the applications presented in [13] and [14], we decomposed the images to be registered into 8 scales each. We then register the coarsest scales with one another, and use the resulting transformation as an initial guess to register the next scales with one another. We then iterate this procedure until the finest scale is reached, at each stage using the transformation parameters obtained in the previous registration as the initial parameters for the current registration. Through several registration experiments in both two and three dimensions, we showed that our multiscale techniques produce significantly more accurate rigid and deformable registration results than standard techniques.

In this paper, we present a novel algorithm for the deformable registration of medical images. Our goal is to improve the accuracy of our previous multiscale deformable registration algorithm by incorporating more information about specific aspects of the images to be registered into the registration process. In particular, the information that we wish to include is the location of known

landmarks, such as corresponding bony structures or other anatomical features, in the images to be registered. To do this, we derive a hybrid image registration algorithm that combines landmark-based registration, deformable splines-based registration, and the multiscale image decomposition.

Our first step is to use the hierarchical multiscale image decomposition of [21] to decompose the images to be registered into *coarse* and *fine* scales. In the coarse scale of each image, which contain its main shapes and general features, we manually match several pairs of corresponding landmarks (e.g. bony structures). We then register the *coarse scales* of the images with one another using a landmark-based registration algorithm. The resulting transformation serves as an *initial guess* for the multiscale deformable registration of the original images with one another (using the multiscale deformable registration algorithm that we presented in [14]). That is, we first use the transformation produced by the coarse scale landmark-based registration as the starting point for a splines-based deformable registration of the coarsest scales with one another. We then use the resulting transformation as the starting point to register the next scales with one another, and iterate the process until the finest scales are reached.

This paper is organized in the following way. In Section 2, we briefly review the image registration problem, discuss landmark and deformable registration techniques, and review the hierarchical multiscale image decomposition of [21]. In Section 3 we develop our hybrid image registration algorithm based on combining the multiscale decomposition with landmark and deformable registration techniques. The accuracy and efficiency of our hybrid technique is demonstrated in Section 4 with several image registration case studies in both two and three dimensions using four-dimensional computed tomography (4D CT) images. In particular, we demonstrate that our hybrid registration algorithm is more accurate than standard deformable registration techniques, standard landmark-based registration techniques, our previous multiscale deformable registration techniques, and combined standard landmark-based and deformable registration techniques. In Section 5, we demonstrate the robustness of our algorithm with respect to the location of landmarks, number of landmarks, and presence of noise. One of the main advantages of our technique over ordinary landmark-based techniques is that our algorithm is not sensitive to the particular location or exact correspondence of the landmarks used in the coarse scale landmark-based phase. Thus our technique significantly reduces the time required for identification of corresponding landmarks compared to standard landmark-based registration techniques. Concluding remarks are given in Section 6.

2 Background

2.1 The registration problem: a mathematical setting

Given two images, a *fixed* and a *moving* image, the classical registration problem is to find an *optimal spatial transformation* that maps the moving image into the fixed image. While this problem is easy to state, it is difficult to solve. The main source of difficulty is that the problem is ill-posed, which means, for example, that the problem may not have a unique solution. Additionally, the notion of optimality may vary for each application. Finally, computation time and data storage constraints place limitations on the complexity of models that can be used for describing the problem.

A two-dimensional gray-scale image f is a mapping which assigns to every point $x \in \Omega \subset \mathbb{R}^2$ a gray value $f(x)$ (called the intensity value of the image at the point x). We will consider images as elements of the space $L^2(\mathbb{R}^2)$. Extensions to color images can be defined, for example, in terms of vector-valued functions $\mathbf{f} = (f_1, f_2, f_3)$ representing the RGB-color scales. For the medical imaging applications that we are interested in, images are in fact given in terms of discrete data, and the

function f must be obtained via interpolation. We will not discuss this construction here, but assume that an interpolation method has been chosen.

As an input one is given two images: the first is defined as the fixed image A and the second as the moving image B . The goal is then to find a transformation ϕ such that the fixed image A and the deformed moving image $B_\phi(x) := B(\phi(x))$ are *similar*. To solve this problem in a mathematical way, the term “similar” needs to be defined in an appropriate fashion. For example, if the images to be registered are taken from different devices, there may not be a correspondence between the intensities A and B_ϕ for an optimal ϕ . Thus the registration problem necessarily involves a discussion of the distance measures, or metrics, used to compare images. Common metrics used in image registration are mean squares, normalized cross correlation, and mutual information [12].

2.2 Landmark-based registration

Landmark-based registration is an image registration technique which is based on a finite set of image features. The problem is to determine the transformation such that for a finite set of features, any feature of the moving image is mapped onto the corresponding features of the fixed image. More precisely, let $F(A, j)$ and $F(B, j)$, $j = 1, \dots, m$ be given features of the fixed (A) and moving (B) images, respectively. The solution ϕ of the registration problem is then a map $\phi : \mathbb{R}^2 \rightarrow \mathbb{R}^2$ (in the two-dimensional case) such that

$$F(A, j) = \phi(F(B, j)), \quad j = 1, \dots, m. \quad (1)$$

For a more general notion of landmark-based registration, we define the following distance measure:

$$D^{LM}(\phi) := \sum_{j=1}^m \|F(A, j) - \phi(F(B, j))\|_l^2, \quad (2)$$

where $\|\cdot\|_l$ denotes a norm on the landmark, or feature, space. For example, if the features are locations of points, then $\|\cdot\|_l = \|\cdot\|_{\mathbb{R}^2}$. We can then restate (1) as the minimization problem in which the solution $\phi : \mathbb{R}^2 \rightarrow \mathbb{R}^2$ of the registration problem is given by:

$$\phi = \underset{\psi: \mathbb{R}^2 \rightarrow \mathbb{R}^2}{\operatorname{argmin}} D^{LM}(\psi). \quad (3)$$

To solve the minimization problem (3), the transformation ϕ is chosen either to be an element of an n -dimensional space spanned, for example, by polynomials, splines, or wavelets, or it is required to be smooth in some sense. In the first case, the features to be mapped are the locations of a number of user-supplied landmarks. Let χ_k , $k = 1, \dots, n$ be the basis functions of the space. Then the minimization of

$$D^{LM}(\phi) := \sum_{j=1}^m \|(A, j) - \phi(F(B, j))\|_l^2$$

can be obtained upon expanding $\phi = (\phi_1, \phi_2)$ in terms of the basis functions χ_k and solving the resulting least squares problems.

Landmark-based registration is simple to implement and the numerical solution requires only solving a linear system of equations. However, the main drawback of the landmark-based approach is that the registration process depends on the location of the landmarks. As the detection and

mathematical characterization of landmarks (for example, anatomical landmarks in medical images) is not fully automated, the landmarks must be user-supplied, and this can be a time-consuming and difficult process, even for a medical expert; see, for example, [16]. Additionally, landmark-based registration does not always result in a physically meaningful registration. In other words, mapping a given set of landmarks from one image to another does not always provide enough constraints to guarantee an acceptable solution. See [12, p. 44], for a simple example of a situation in which landmark-based registration fails to produce meaningful results.

2.3 Deformable registration

An alternative approach is to use methods that are based on the minimization of some distance measure, or metric, D . The transformation ϕ is restricted to some parameterized space, and the registration can be obtained by minimizing the distance D over the parameterized space. Given a metric D , a fixed image A , and a moving image B , optimal parametric registration is the problem of finding a transformation ϕ in some pre-specified parameterizable space such that $D(A, B(\phi))$ is minimized.

Examples of commonly used parameterizable spaces in optimal parametric image registration are rigid, affine, polynomial, and spline spaces. Spline-based transformation models are among the most common and important transformation models used in non-rigid registration problems [6]. Spline-based registration algorithms use *control points* in the fixed and moving images and a spline function to define transformations away from these points. The two main spline models used in image registration problems are *thin-plate splines* and *B-splines*. Thin-plate splines have the property that each control point has a *global* influence on the transformation. That is, if the position of one control point is perturbed, then all other points in the image are perturbed as well. This can be a disadvantage because it limits the ability of the transformation model to model localized deformations. In addition, the computation time required for a thin-plate spline-based registration algorithm increases significantly as the number of control points increases. See [3] for an overview of thin-plate splines.

In contrast, B-splines are defined *locally* in the neighborhood of each control point. Thus perturbing the position of one control point affects the transformation only in a neighborhood of that point. As a result, B-spline-based registration techniques are more computationally efficient than thin-plate splines, especially for a large number of control points. See [7] for a detailed description of B-splines transformation models. The displacement of a node $\alpha_{i,j}$ is specified by a vector $\mathbf{x}_{i,j}$, and the displacement vectors of a collection of nodes characterize the tissue deformation. The displacement at a particular location on the image is deduced by fitting a polynomial expressed using the basis spline (B-spline) to the grid nodes. The number of control points used determines the number of degrees of freedom of the transformation model, and hence, the computational complexity of the B-splines algorithm.

2.4 The multiscale decomposition

We briefly review the hierarchical (BV, L^2) decompositions of [21]. This multiscale decomposition will provide a hierarchical expansion of an image that separates the essential features of the image (such as large shapes and edges) from the fine scales of the image (such as details and noise). The decomposition is hierarchical in the sense that it will produce a series of expansions of the image that resolve increasingly finer scales, and hence include increasing levels of detail.

We shall use in the decomposition the space of functions of bounded variation, BV , defined by:

$$BV = \left\{ f \mid \|f\|_{BV} := \sup_{h \neq 0} |h|^{-1} \|f(\cdot + h) - f(\cdot)\|_{L^1} < \infty \right\}.$$

Define the J -functional $J(f, \lambda)$ as follows:

$$J(f, \lambda) := \inf_{u+v=f} \lambda \|v\|_{L^2}^2 + \|u\|_{BV}, \quad (4)$$

where $\lambda > 0$ is a scaling parameter (to be determined by the user) that separates the L^2 and BV terms. This functional $J(f, \lambda)$ was introduced in the context of image processing by Rudin, Osher, and Fatemi [17]. Let $[u_\lambda, v_\lambda]$ denote the minimizer of $J(f, \lambda)$. This minimization is interpreted as a decomposition $f = u_\lambda + v_\lambda$, where u_λ extracts the edges of f and v_λ extracts the textures of f . However, this interpretation depends on the scale λ , since texture at scale λ consists of edges when viewed under a refined scale (2λ , for example). We refer to $v_\lambda = f - u_\lambda$ as the *residual* of the decomposition. Upon decomposing $f = u_\lambda + v_\lambda$, we proceed to decompose v_λ as follows:

$$v_\lambda = u_{2\lambda} + v_{2\lambda},$$

where

$$[u_{2\lambda}, v_{2\lambda}] = \operatorname{arginf}_{u+v=v_\lambda} J(v_\lambda, 2\lambda).$$

Thus we obtain a two-scale representation of f given by $f \cong u_\lambda + u_{2\lambda}$, where now $v_{2\lambda} = f - (u_\lambda + u_{2\lambda})$ is the residual of the decomposition. This process can be repeated. Denoting the initial scale $\lambda = \lambda_0$, we have, after k steps,

$$f = \sum_{i=0}^k u_i + v_k, \quad (5)$$

where

$$[u_{j+1}, v_{j+1}] = \operatorname{arginf}_{u+v=v_j} J(v_j, \lambda_0 2^{j+1}), \quad j = 0, 1, \dots$$

We remark that this decomposition, although presented originally for two-dimensional images, has a similar form in any number of dimensions. The details of the implementation in two and three dimensions is given in Appendix A.

For ease of notation, given an image f , we let $C_k(f)$ denote the k -th scale of the image f , $k = 1, \dots, m$:

$$C_k(f) = \sum_{i=0}^{k-1} u_i(f). \quad (6)$$

In particular, note that $C_1(f)$ is the *coarsest scale* of the image f , containing only the most general features and main shapes of the image. As k increases, the components $C_k(f)$ contain increasing amounts of detail (and noise), and the last component $C_m(f)$ is a close approximation to the original image.

3 Hybrid Landmark and deformable registration using the multiscale decomposition

In our previous work [13] and [14], we used the multiscale decomposition of [21] reviewed in Section 2.4 to develop multiscale rigid and deformable registration algorithms for registration of images that contain high levels of noise. We demonstrated that our multiscale registration techniques enable successful registration of images that contain noise levels significantly greater than the levels at which ordinary rigid and deformable registration techniques fail. In this paper, we extend our previous work on multiscale registration to a new registration technique based on combining landmark and deformable registration methods with the multiscale decomposition. We shall refer to this technique as a *hybrid multiscale landmark and deformable registration* algorithm. The main idea behind our new technique is that by incorporating landmark registration, we can use known information about correspondence of bony structures or other anatomical structures to improve the accuracy and efficiency of the registration procedure. This is done by providing a more accurate initial guess to the deformable registration algorithm.

For the general setup, consider two images A (the fixed image) and B (the moving image), and suppose that we want to register image B with image A . We first apply the multiscale hierarchical decomposition to both images. Let m denote the number of hierarchical steps used for the multiscale decompositions.

Upon decomposing both of the images to be registered, we use the coarse scales of each image to identify several corresponding anatomical landmarks in each image. To identify landmarks, we must choose anatomical features of the two coarse scales that we know correspond to one another. For example, bony structures (such as the spine or ribs in images of the torso) in one image must correspond to bony structures in another image.

We then use the identified landmarks to perform a landmark-based registration of the coarse scales of each image. Since the coarse scales of the image contain only the main shapes and general features, identification of corresponding landmarks in the coarse scales is a relatively easy task, and can potentially be implemented automatically (for example, using the methods of [20]). Although in many cases corresponding features can also be easily identified in the original images, the initial coarse scale landmark-based registration serves as a computationally efficient initial registration whose final registration parameters are used as the initial registration parameters for a deformable B-splines registration, or a multiscale deformable B-splines registration, of the original images. More precisely, the output of the landmark registration algorithm is the transformation ϕ_{landmark} that optimally brings the first coarse scale of the moving image, $C_1(B)$, into spatial alignment with the first coarse scale of the fixed image, $C_1(A)$, such that corresponding landmarks are mapped to one another. This coarse-scale transformation is thus a reasonable approximation for the optimal spatial transformation between the original images A and B , and is used as an initial guess, or starting point, for the deformable registration of the original images with one another.

This deformable registration can be implemented in one of two ways. First, the original images A and B can be registered with one another using a standard B-splines deformable registration algorithm with the output of the landmark registration ϕ_{landmark} as the starting point for the algorithm.

Alternatively, the images can be registered with one another using the *multiscale* deformable registration algorithm that we proposed in [14]. To implement this algorithm, the scales of each image are registered iteratively with one another, at each stage using the transformation parameters from the previous registration as the input parameters for the current registration. To begin the

multiscale deformable registration, the first coarse scales $C_1(A)$ and $C_1(B)$ are registered with one another using a deformable B-splines algorithm with the output of the landmark registration ϕ_{landmark} as the starting point. The resulting transformation is then used as the starting point to register the next scales $C_2(A)$ and $C_2(B)$ with one another, and the process is iterated until the final scales $C_m(A)$ and $C_m(B)$ are reached. The multiscale deformable registration algorithm is illustrated schematically in Figure 1.

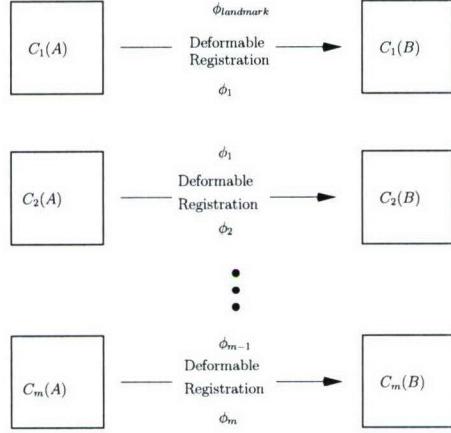
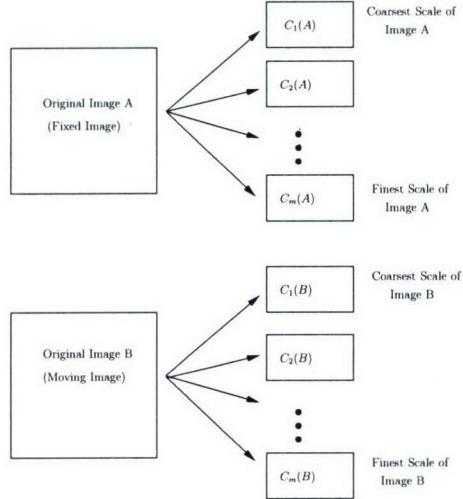


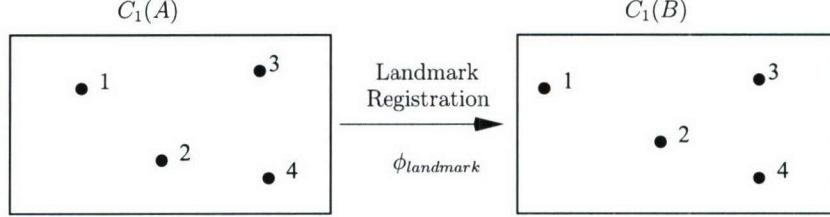
Figure 1: A schematic representation of the iterated multiscale deformable registration algorithm.

To summarize, the implementation of the hybrid multiscale landmark and deformable registration algorithm is as follows:

1. Decompose the fixed and moving images to be registered using the hierarchical multiscale image decomposition of [21] reviewed in Section 2.4.



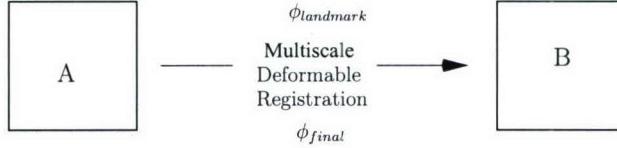
2. Identify pairs of corresponding landmarks in the coarse scales of the fixed and moving images, and use a standard landmark-based registration algorithm to register the coarse scales with one another.



3. (a) Use a B-splines deformable registration algorithm to register the original moving image with the original fixed image using the transformation computed via the coarse-scale landmark-based registration in Step 2 as an *initial guess* for the optimal transformation that brings the original images into alignment with one another.



- (b) Use a multiscale deformable registration algorithm [14] to register the original moving image with the original fixed image using the transformation computed via the coarse-scale landmark-based registration in Step 2 as an *initial guess*. In the multiscale algorithm, the scales of each image are registered iteratively with one another, as illustrated in Figure 1.



The development of our hybrid multiscale technique was motivated by the observation that the correspondence between some of the regions in the moving image, such as bony structures, can be easily identified visually, while the correspondence between other regions, such as those that contain tissue deformation, breathing movement, or lack of distinct image features, is less obvious. Thus mapping of the two types of regions should be approached differently, rather than equally as implemented in ordinary image registration algorithms. Our proposed technique combines landmark registration and deformable registration so that prior knowledge about the correspondence between the images, such as visual identification of corresponding landmarks, can be incorporated into the deformable registration process. The use of the coarse scales (which contain only the main shapes and general features of the images) obtained via the hierarchical multiscale image decomposition of [21] enables quick and easy identification of corresponding landmarks, and the deformable registration component of the algorithm fine-tunes the registration result produced by

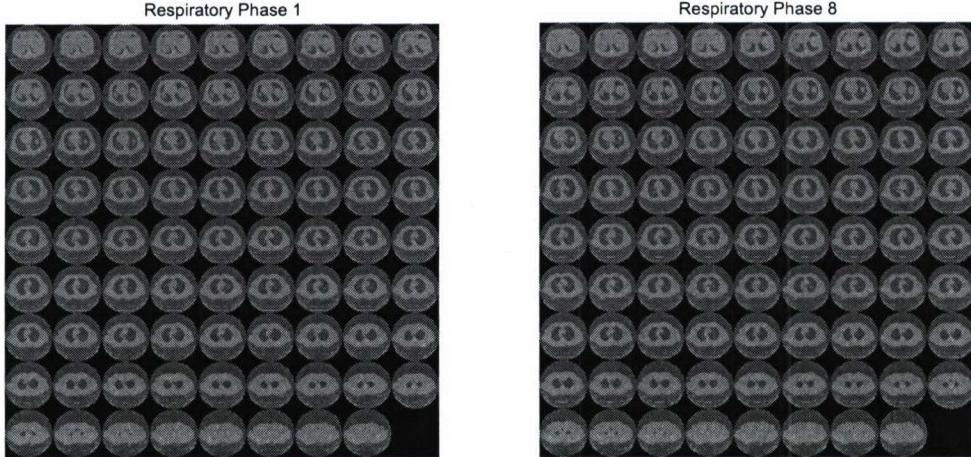


Figure 2: The CT image slices (80 per phase) corresponding to the first and eighth breathing phases of the respiratory cycle.

the initial coarse-scale landmark registration. The key feature of our hybrid algorithm is that the transformation produced by the coarse scale landmark-based registration algorithm is used as the starting point, or initial guess, for the deformable registration component of the algorithm. The initial coarse scale registration transformation is an approximation for the optimal transformation that brings the original images into alignment with one another, so using the coarse scale landmark-based transformation as a starting point for deformable registration reduces the computation time needed for deformable registration and should prevent the deformable registration algorithm from converging to non-meaningful results.

4 Results and discussion

In this section, we demonstrate the accuracy, efficiency, and robustness of the image registration algorithm presented in Section 3 with image registration experiments using both clinical and synthetic deformations. The images used in this section were acquired with a GE Discover-ST Scanner (GE Medical Systems, Milwaukee, WI) at the Stanford University Medical Center. We obtained eight sets (*phase bins*) of images, each consisting of 80 two-dimensional computed tomography (CT) images (*slices*) of the lungs. Each phase corresponds to a different breathing phase of the respiratory cycle. Each 2D image slice contains 512×512 voxels, and the slice thickness for each phase is 2.5-mm, and the eight breathing phases recorded contain approximately 400 MB of data in DICOM image format.

In Figure 2, we illustrate the CT image slices corresponding to the first and eight breathing phases of the respiratory cycle.

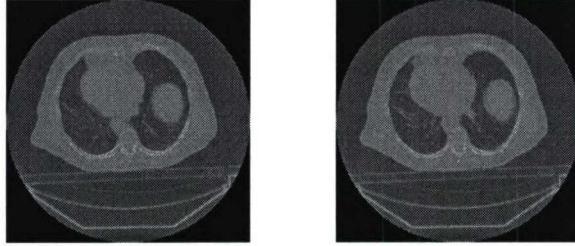


Figure 3: Two corresponding sample slices from two breathing phases of the same patient. The image on the left (denoted $P_{1,16}$) is from the first breathing phase, and the image on the right (denoted $P_{8,16}$) is the corresponding slice from the eighth breathing phase.

4.1 2D registration of respiratory phases

We consider two corresponding sample slices (slice 16) from the first and eighth phases of the CT data set (shown in Figure 3). We register these images using both ordinary deformable registration and the multiscale hybrid registration algorithm. We let $P_{1,16}$ denote the fixed image (on the left), and $P_{8,16}$ denote the moving image (on the right).

To implement the multiscale hybrid algorithm, we first decompose each of the images into $m = 2$ hierarchical scales. The first coarse scales of each of the slices are illustrated in Figure 4.

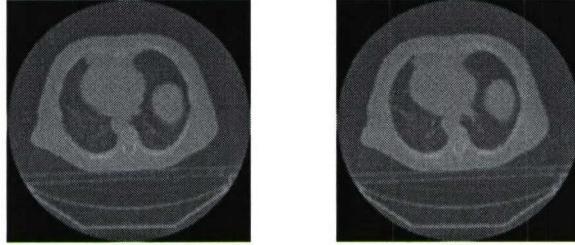


Figure 4: The first coarse scales of $P_{1,16}$ and $P_{8,16}$ shown in Figure 3.

To implement the landmark-based registration portion of the hybrid algorithm, we identify four corresponding pairs of landmarks with one another, as shown in Figure 5.

Next, we use a landmark-based registration algorithm to register the images using the features identified in Figure 5. The result of this registration process is a transformation ϕ which optimally maps the first coarse scale of the moving image, $C_1(P_{8,16})$, to the first coarse scale of the fixed image, $C_1(P_{1,16})$ such that the corresponding landmarks identified in Figure 5 are mapped to one another. Finally, we apply the transformation ϕ to the original moving image $P_{8,16}$, and register the resulting transformed image with the fixed image $P_{1,16}$ using a B-splines deformable registration algorithm.

In Figure 6, we illustrate the difference between the slices before registration and after registra-

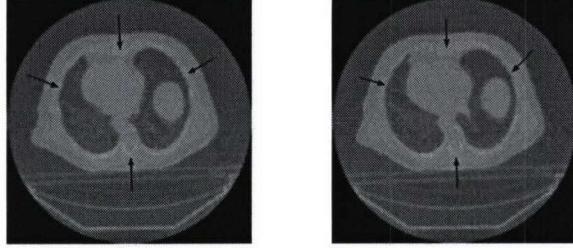


Figure 5: The four landmarks that are identified on the coarse scales of $P_{1,16}$ (left) and $P_{8,16}$ (right).

tion performed using the multiscale hybrid registration technique. For comparison, we also illustrate the difference between the slices after standard B-splines deformable registration (without using the multiscale decomposition) and after landmark registration. The visual results presented in Figure 6 clearly demonstrate the accuracy of the multiscale hybrid registration technique. The hybrid algorithm successfully recovers the deformation between the two images, and is, indeed, more accurate than both ordinary deformable registration and landmark registration. Moreover, the hybrid algorithm is computationally more efficient than ordinary deformable registration techniques. Working on a Dell Dimension 8400 Intel Pentium 4 CPU (3.40 GHz, 2.00 GB of RAM), ordinary deformable registration of the images shown in Figure 3 requires approximately 115 seconds. The total time required for the hybrid registration algorithm (including the multiscale decomposition of both images, landmark registration of the coarse scales, and the final deformable registration) is approximately 72 seconds.

To further quantitatively evaluate the accuracy of the hybrid registration technique, we use the multiscale hybrid algorithm to register all 80 slices of the first breathing phase with the corresponding slices of the eighth breathing phase, and for each registration we compute the correlation coefficient ρ between the slices before and after registration. The correlation coefficient $\rho(A, B)$ between two images A and B is given by:

$$\rho(A, B) = \frac{\sum_m \sum_n (A_{mn} - \bar{A})(B_{mn} - \bar{B})}{\sqrt{\sum_m \sum_n (A_{mn} - \bar{A})^2 (B_{mn} - \bar{B})^2}},$$

where A and B are $m \times n$ two-dimensional images and \bar{A} and \bar{B} represent the mean value of the elements of A and B , respectively. ($\rho = 0$ indicates a low degree of matching between the images, while $\rho = 1$ indicates that the images are identical). In Figure 7, we plot the correlation coefficients between the slices before registration, after hybrid multiscale registration, after multiscale deformable registration, after ordinary deformable registration, and after landmark registration.

The correlation coefficients plot in Figure 7 quantitatively confirms the accuracy of the multiscale hybrid registration technique. For all 80 CT lung slices considered, the hybrid technique significantly improves the correlation coefficient between corresponding images, indicating that the algorithm successfully recovers the deformation between the breathing phases. Moreover, the hybrid multiscale technique is more accurate than all other registration methods considered. Finally, in Figure 8, we compare the correlation coefficients obtained with hybrid registration (coarse scale landmark

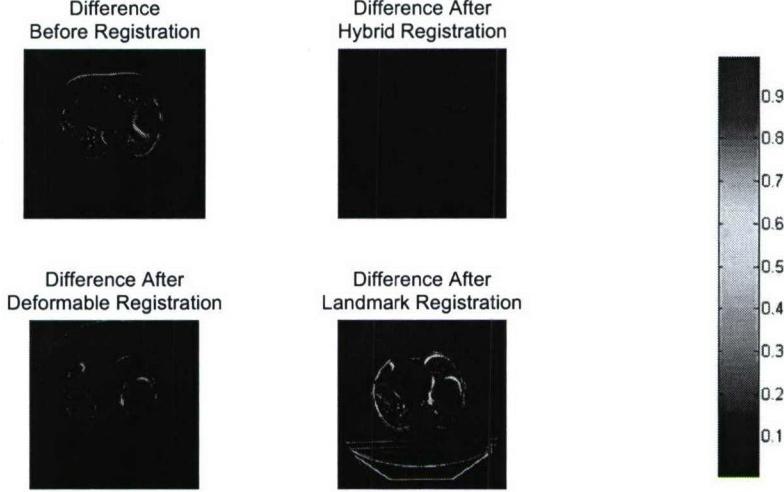


Figure 6: The difference between the CT respiratory phase slices before registration (upper left), after hybrid registration (upper right), after deformable registration (lower left), and after landmark registration (lower right).

registration followed by standard deformable registration) and hybrid multiscale registration (coarse scale landmark registration followed by multiscale deformable registration). The results in Figure 8 demonstrate that hybrid multiscale registration improves the accuracy of hybrid registration. For all 80 slices, the correlation coefficient between the images is between 0.98 and 1.0 after hybrid multiscale registration.

4.2 Large deformation registration

In this section, we demonstrate the accuracy of the hybrid technique for registration images that contain large localized deformations. We apply a large splines deformation to the fixed image $P_{1,16}$ (slice 16, first breathing phase). The corresponding deformation field and the fixed and deformed images are illustrated in Figure 9.

As in Section 4.1, we first decompose the images to be registered into $m = 2$ hierarchical scales. Using the first coarse scale of each image, we identify four corresponding pairs of bony structure landmarks, and use landmark-based registration to register the coarse scales of the images with one another. We then apply the resulting transformation to the original deformed image, and use a B-splines deformable registration algorithm to complete the registration process. In Figure 10, we compare the difference between the images before registration, after registration performed using the hybrid registration technique, and (for reference) after ordinary deformable registration and landmark registration.

The results presented in Figure 10 demonstrate that the hybrid registration technique successfully registers images that contain large localized deformations. For such images, the multiscale

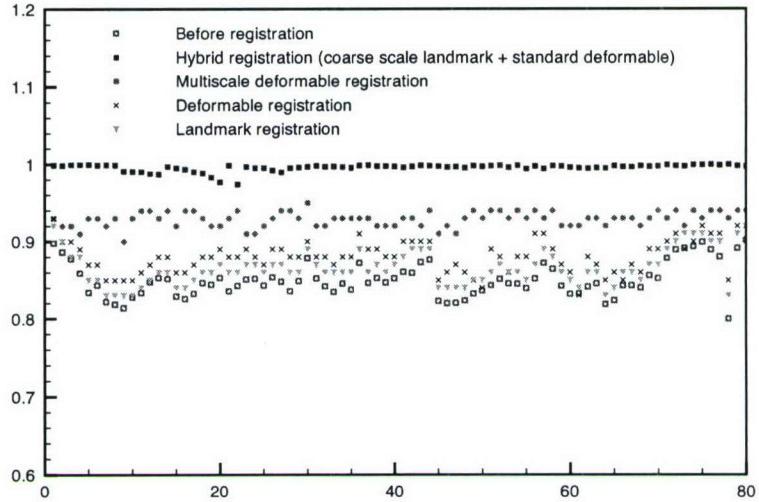


Figure 7: The correlation coefficients between the first and eight breathing phase image slices before registration, after hybrid multiscale registration, after multiscale deformable registration, after ordinary deformable registration, and after landmark registration.

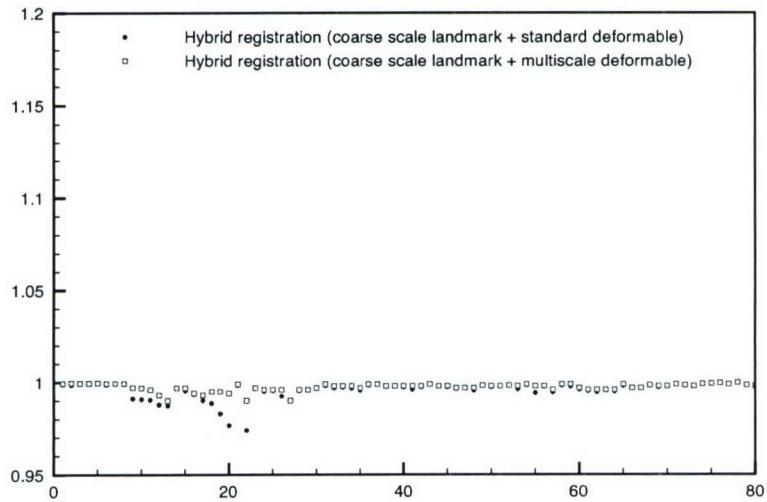


Figure 8: The correlation coefficients between the first and eight breathing phase image slices after hybrid registration and after hybrid multiscale registration.



Figure 9: The deformation field (left) corresponding to the deformation transformation between the fixed (center) and deformed (right) images to be registered with one another.

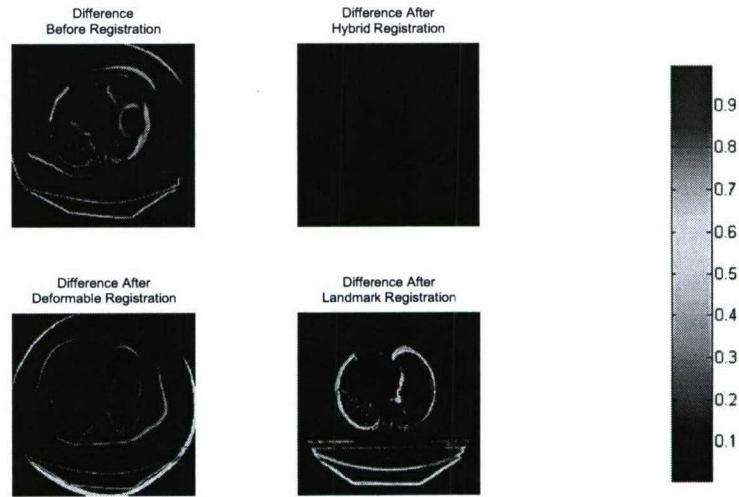


Figure 10: The difference between the fixed and deformed images with a large deformation (Figure 9) before registration (upper left), after hybrid registration (upper right), after deformable registration (lower left), and after landmark registration (lower right).

Registration Method	Hybrid	Multiscale	Deformable	Landmark
ρ after registration	0.97	0.93	0.90	0.84

Table 1: The correlation coefficients between the fixed and deformed images before registration, after hybrid multiscale registration, after multiscale deformable registration, and after landmark registration for the large deformation test case shown in Figure 9. Before registration $\rho = 0.70$.

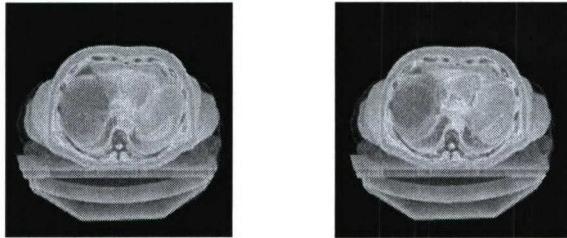


Figure 11: The 3D CT volumes obtained upon combining the 80 2D CT phase-binned images corresponding to the first (left) and eighth (right) respiratory phases.

hybrid registration technique is significantly more accurate than both ordinary deformable registration and ordinary landmark registration. The improvement of the multiscale hybrid technique over ordinary deformable registration is much more pronounced in this case, indicating that the multiscale hybrid registration algorithm is particularly well-suited for registration problems involving large localized deformations. In Table 1, we compare the correlation coefficients between the fixed and deformed images before registration, after hybrid multiscale registration, after multiscale deformable registration, and after landmark registration.

4.3 3D Registration of respiratory phases

Although the two-dimensional slice-by-slice registration results presented in Figures 7 and 8 demonstrate that our registration algorithm registers corresponding slices from different volumes with one another, it is more physically meaningful to perform a three-dimensional registration of the volumes with one another, as corresponding slices in different phases may represent different anatomy components. Thus, in this section, we use our multiscale hybrid registration technique to accurately register three-dimensional images. We combine the 80 two-dimensional images from respiratory phases 1 and 8 (shown in Figure 2) to obtain two three-dimensional CT images, as illustrated in Figure 11.

In our work [14] on multiscale deformable registration, we extended the hierarchical multiscale image decomposition to three-dimensional images. Upon decomposing each 3D image, we identify corresponding landmarks in each image. For the purposes of this paper, we identify the landmarks slice-by-slice (as in the 2D case), and refer the interested reader to [20] for an algorithm to automatically identify homologous regions in each slice. We then register the 3D coarse scales with one another using the identified landmarks, and apply the resulting transformation to the original moving image. Finally, we use a deformable B-splines registration algorithm to register the transformed

moving image with the fixed image (as in the 2D case). In Figure 12, we compare the intensity difference between two sample slices before and after multiscale hybrid registration. A comparison of the intensity differences before and after registration demonstrates that the multiscale hybrid registration method indeed recovers the difference between the two images. Similar results were obtained with all other slices. The correlation between the images is 0.83 before registration and 0.97 after registration. The total required computation time for both the 3D multiscale decomposition and multiscale hybrid registration algorithm is on the order of approximately 30-50 minutes, depending on the data set; this particular example required 42 minutes. We would like to note that as an alternative approach to 3D registration, one can register the 2D slices in parallel. While this will not take into account all possible 3-dimensional deformations, the algorithm will be significantly faster (of the order of several minutes).

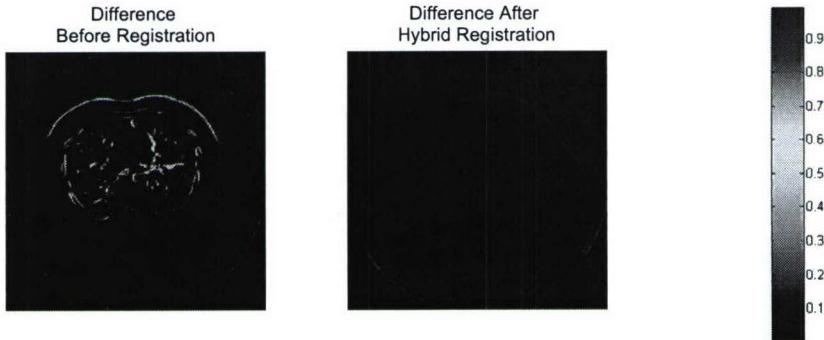


Figure 12: The difference between two corresponding slices of the 3D CT images before and after multiscale hybrid registration.

5 Robustness of the multiscale hybrid registration algorithm

In this section, we demonstrate the robustness of the multiscale hybrid registration algorithm with several image registration experiments. All image registration experiments in this section are performed using the fixed and deformed images shown in Figure 9. The fixed image in each experiment is the image $P_{1,16}$ (first respiratory phase, slice 16), and the moving image is the image obtained upon applying a large B-splines deformation to the fixed image. We shall consider robustness with respect to the location of the landmarks, number of landmarks, and several types of noise.

5.1 Location of landmarks

We start with a study of the robustness of the multiscale hybrid registration technique with respect to the location of the bony structure landmarks identified using the coarse scale representations of the images to be registered. We perform 20 image registration experiments in which we randomly perturb the position of the coarse-scale landmarks in the fixed image only by 10-20 mm in each

ρ after registration	0.96	0.97	0.98	0.99
Number of trials (out of 20)	4	8	6	2

Table 2: The correlation coefficients between the fixed and deformed images after multiscale hybrid registration. Each trial number represents a different perturbation of the landmarks in the fixed image. Before the registration $\rho = 0.70$. For the exact location of the landmarks: $\rho = 0.96$.

trial (the image sizes are all 512 pixels \times 512 pixels). In each trial, we place the landmarks in the moving image in exactly the same positions as those illustrated in Figure 5, but we vary the location of each corresponding landmark in the fixed image randomly by 10-20 mm, as illustrated in Figure 13. The moving image landmarks are randomly placed in the blue circles shown in Figure 13.

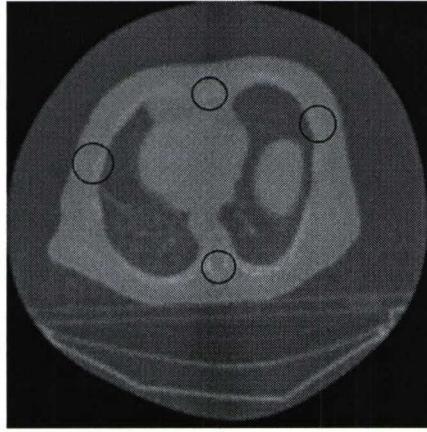


Figure 13: The regions for the perturbed locations of the coarse-scale landmarks in the moving images. In each trial, the landmarks in the moving image are randomly placed in the blue circles.

These experiments are designed to determine whether or not the accuracy of the hybrid algorithm is dependent on precise matching between the landmarks. In Table 2, we present the correlation coefficients between the images after registration. The correlation coefficient in all trials is in the range $\rho = 0.96\text{--}0.98$. Before registration, $\rho = 0.70$. For reference, we also note that $\rho = 0.96$ for the exact location of the landmarks.

The results presented in Table 2 demonstrate that the multiscale hybrid registration technique is computationally robust with respect to the location of the landmarks used in the coarse-scale landmark registration phase of the algorithm. Although the precise location of the landmarks differs in each of the 20 trials, the final registration results are essentially the same. This is because the coarse-scale landmark registration is only the first step in the hybrid registration algorithm; the deformable registration step fine-tunes the result of the landmark registration, thus correcting any anomalies introduced in the landmark selection process. This robustness indicates that the hybrid algorithm is particularly well-suited for clinical applications. Although the algorithm does require the identification of corresponding landmarks in the coarse-scale images, the algorithm is

Number of landmarks	2	3	4	5	6	7
ρ after registration	0.93	0.94	0.96	0.95	0.94	0.95
Total computation time (seconds)	99	89	75	77	85	86

Table 3: The correlation coefficients between the fixed and deformed images after hybrid registration and total computation time required for hybrid registration for image registration experiments using a varying number of landmarks.

not sensitive to the exact location of the landmarks.

This robustness to the location of the landmarks is one of the most notable features of our hybrid multiscale landmark and deformable registration algorithm. Most landmark-based registration algorithms require precise identification of corresponding landmarks and are thus tedious, time-consuming, and difficult to implement, even for a medical expert. However, since the accuracy of our hybrid algorithm is not dependent on precise correspondence of the landmarks in the fixed and moving images, the identification of the landmarks can be completed quickly or implemented automatically.

5.2 Number of landmarks

In this section, we consider the effect of the number of pairs of landmarks used in the landmark-based registration step on the accuracy and efficiency of the multiscale hybrid registration algorithm. In Table 3, we present the correlation coefficients between the images after registration and total registration time (including the multiscale decomposition, landmark-based registration of the coarse scales, and deformable registration) for image registration experiments using a varying number of landmarks. The locations of the landmarks are illustrated in Figure 14. The results presented in Table 3 indicate that accurate registration can be obtained using 4 pairs of corresponding landmarks.

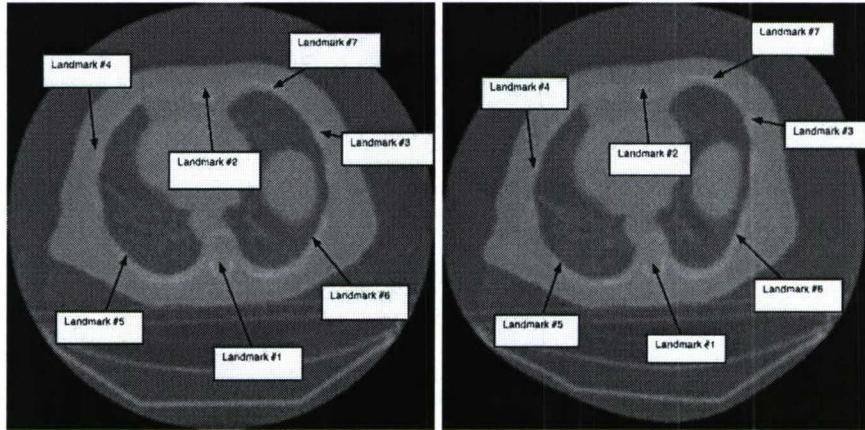


Figure 14: The locations of the landmarks used in the experiments presented in Table 3.

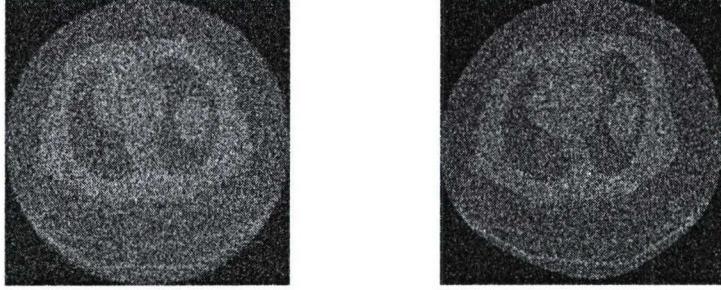


Figure 15: The fixed and deformed images with Gaussian noise of mean 0 and variance 0.1.

5.3 Noise

In this section, we demonstrate the robustness of the multiscale hybrid registration algorithm with respect to the presence of noise in the images to be registered. We consider three types of noise models: Gaussian, multiplicative, and impulse noise. For each type of noise, we let $f(x)$ denote the observed (noisy) image and let $s(x)$ denote the true image. See [13] and [14] for a detailed description of the problem of image registration in the presence of noise. Gaussian noise is modeled by the following equation:

$$f(x) = s(x) + n(x), \quad (7)$$

where $n(x) \sim N(\mu, v)$ is uniformly distributed random noise with mean μ and variance v . In Figure 15, we illustrate the fixed and moving images with added Gaussian noise of mean 0 and variance 0.1.

Multiplicative noise is defined by the following model:

$$f(x) = s(x) + \eta(0, \delta) \cdot s(x), \quad (8)$$

where $\eta(0, \delta)$ is uniformly distributed random noise of mean 0 and variance δ .

Impulse noise, or salt and pepper noise, is typically defined by the following model:

$$f(x) = \begin{cases} s(x), & \text{with probability } 1 - \delta, \\ \eta(x), & \text{with probability } \delta, \end{cases} \quad (9)$$

where $\eta(x)$ is an identically distributed, independent random process which sets corrupted pixels alternatively to zero (black) or one (white); unaffected pixels remain unchanged.

In Table 4, we present the correlation coefficients between the noisy images before registration, after ordinary deformable registration, and after hybrid registration. The results are shown for all three noise models.

The registration results presented in Tables 4 clearly demonstrate that our hybrid multiscale algorithm is particularly robust with respect to the presence of noise in the images to be registered. This robustness to noise is one of the main strengths of the hybrid algorithm. We have demonstrated in [13] and [14] that ordinary registration techniques fail to produce accurate results when the images to be registered are noisy. The hybrid multiscale method presented here significantly improves both ordinary registration techniques and the multiscale registration techniques that we presented in [14] for the problem of noisy image registration.

Registration Method	Gaussian $N(0, 0.1)$		Multiplicative $\eta(0, 0.2)$		Impulse $\delta = 0.2$	
	ρ	T	ρ	T	ρ	T
Before registration	0.18	—	0.35	—	0.22	—
Deformable	0.20	304	0.46	295	0.25	333
Multiscale Deformable	0.89	92	0.90	89	0.90	105
Hybrid	0.91	81	0.93	76	0.93	82

Table 4: The correlation coefficients between the noisy fixed and moving images shown in Figure 15 before registration, after ordinary deformable registration, and after multiscale hybrid registration. ρ is the correlation coefficient after registration. T denotes the registration time in seconds.

6 Conclusions

In this paper, we extended our previous works [13] and [14] on multiscale image registration for accurate registration of noisy images. We proposed a hybrid two-step image registration technique combining the hierarchical multiscale image decomposition of [21] with landmark registration methods, deformable registration methods, and our previous multiscale deformable registration methods. Our hybrid technique allows the practitioner to incorporate *a priori* knowledge of corresponding bony or other anatomical structures into the registration process by using a landmark registration algorithm to register the coarse scales of the images to be registered with one another. The transformation produced by this coarse scale landmark registration is then used as the starting point for a B-splines deformable (standard or multiscale) registration of the original images with another. Using the coarse scale landmark registration as an initial guess for the deformable phase of the algorithm is the main new feature of our hybrid algorithm. The initial guess provided by the landmark-based registration significantly increases the accuracy of the registration process and decreases the computation time required for convergence of the algorithm.

With extensive registration experiments in both two and three dimensions using 4D CT images, we have shown that the hybrid technique is more accurate than ordinary landmark registration, ordinary deformable registration, and multiscale deformable registration, especially in the case of registering images that contain large localized deformations and/or noise. Additionally, we have shown that the hybrid method is particularly robust to the location and number of the landmarks used in the coarse scale landmark-based phase of the algorithm, as well as to the presence of noise in the images to be registered. One of the main features of our hybrid technique over standard landmark-based techniques is its robustness with respect to the location of corresponding landmarks. In Table 2, we demonstrated that accurate registration results are obtained even when the location of the landmarks is varied randomly by 10-20 mm. Standard landmark-based techniques, however, are dependent on exact spatial correspondence of landmarks, and are thus tedious and time-consuming to implement. Since the accuracy of our technique is not dependent on exact correspondence of landmarks, our algorithm significantly reduces the time required for the landmark selection portion of such algorithms.

In future research, we will study the problem of auto-detection of coarse scale landmarks for the landmark-based portion of the algorithm. We will also work on embedding the hybrid multiscale landmark and deformable registration ideas presented in this paper into other registration algorithms and applications in which they can be useful.

Acknowledgment: The work of D. Levy was supported in part by the National Science Foundation.

tion under Career Grant No. DMS-0133511. The work of L. Xing was supported in part by the Department of Defense under Grant No. PC040282 and the National Cancer Institute under Grant No. 5R01 CA98523-01.

A Implementation of the multiscale decomposition

We provide the details for the implementation of the hierarchical multiscale image decomposition described in [21]. To begin the decomposition, an initial scale λ_0 must be chosen. For the applications presented in this paper, we will use $\lambda_0 = 0.01$ and $\lambda_j = \lambda_0 2^j$. Due to the singularity when $|\nabla u_\lambda| = 0$, we replace $J(f, \lambda)$ by the regularized functional

$$J^\epsilon(f, \lambda) := \inf_{u+v=f} \left\{ \lambda ||v||_{L^2}^2 + \int_{\Omega} \sqrt{\epsilon^2 + |\nabla u|^2} dx dy \right\}, \quad (10)$$

and at each step, we find the minimizer u_λ of J^ϵ . The Euler-Lagrange equation for $J^\epsilon(f, \lambda)$ is

$$u_\lambda - \frac{1}{2\lambda} \operatorname{div} \left(\frac{\nabla u_\lambda}{\sqrt{\epsilon^2 + |\nabla u_\lambda|^2}} \right) = f \quad \text{in } \Omega,$$

with the Neumann boundary conditions:

$$\frac{\partial u_\lambda}{\partial n} \Big|_{\partial\Omega} = 0, \quad (11)$$

where $\partial\Omega$ is the boundary of the domain Ω and n is the unit outward normal. We thus obtain an expansion $f \sim \sum_{k=0}^m u_k$, where the u_k are constructed as approximate solutions of the following elliptic PDE:

$$u_{k+1} - \frac{1}{2\lambda_{k+1}} \operatorname{div} \left(\frac{\nabla u_{k+1}}{\sqrt{\epsilon^2 + |\nabla u_{k+1}|^2}} \right) = -\frac{1}{2\lambda_k} \operatorname{div} \left(\frac{\nabla u_k}{\sqrt{\epsilon^2 + |\nabla u_k|^2}} \right). \quad (12)$$

To numerically implement the method, we cover the domain Ω with a grid ($x_i := ih, y_j := jh$), and discretize the elliptic PDE of Eq. (12) as follows:

$$\begin{aligned} u_{i,j} &= f_{i,j} + \frac{1}{2h^2} \left[\frac{u_{i+1,j} - u_{i,j}}{\sqrt{\epsilon^2 + (D_{+x}u_{i,j})^2 + (D_{0y}u_{i,j})^2}} - \frac{u_{i,j} - u_{i-1,j}}{\sqrt{\epsilon^2 + (D_{-x}u_{i,j})^2 + (D_{0y}u_{i-1,j})^2}} \right] \\ &\quad + \frac{1}{2h^2} \left[\frac{u_{i,j+1} - u_{i,j}}{\sqrt{\epsilon^2 + (D_{0x}u_{i,j})^2 + (D_{+y}u_{i,j})^2}} - \frac{u_{i,j} - u_{i,j-1}}{\sqrt{\epsilon^2 + (D_{0x}u_{i,j-1})^2 + (D_{-y}u_{i,j})^2}} \right], \end{aligned} \quad (13)$$

where D_+ , D_- , and D_0 denote the forward, backward, and centered divided differences, respectively. Upon solving (13), we obtain:

$$\begin{aligned}
u_{i,j}^{n+1} = & f_{i,j} + \frac{1}{2h^2} \left[\frac{u_{i+1,j}^n - u_{i,j}^{n+1}}{\sqrt{\epsilon^2 + (D_{+x}u_{i,j}^n)^2 + (D_{0y}u_{i,j}^n)^2}} - \frac{u_{i,j}^{n+1} - u_{i-1,j}^n}{\sqrt{\epsilon^2 + (D_{-x}u_{i,j}^n)^2 + (D_{0y}u_{i-1,j}^n)^2}} \right] \\
& + \frac{1}{2h^2} \left[\frac{u_{i,j+1}^n - u_{i,j}^{n+1}}{\sqrt{\epsilon^2 + (D_{0x}u_{i,j}^n)^2 + (D_{+y}u_{i,j}^n)^2}} - \frac{u_{i,j}^{n+1} - u_{i,j-1}^n}{\sqrt{\epsilon^2 + (D_{0x}u_{i,j-1}^n)^2 + (D_{-y}u_{i,j}^n)^2}} \right]. \quad (14)
\end{aligned}$$

To satisfy the Neumann boundary conditions (11), we first reflect f outside Ω by adding grid lines on all sides of Ω . As the initial condition, we set $u_{i,j}^0 = f_{i,j}$. We iterate this numerical scheme for $n = 0, 1, \dots, N$ until $\|u^{n_\infty} - u^{n_\infty-1}\|$ is sufficiently small so that $u_{i,j}^{n_\infty}$ is an accurate approximation of the fixed point steady solution u_λ .

Finally, we denote the final solution $u_\lambda := \{u_{i,j}^{n_\infty}\}_{i,j}$. The full hierarchical multiscale decomposition is obtained by reiterating this process.

To implement the iterated multiscale decomposition in three dimensions, we cover the image domain Ω with a grid ($x_i := ih, y_j := jh, z_k := kh$), and let D_+ , D_- , and D_0 denote the forward, backward, and centered divided differences, respectively. Then the 3D extension of the iterated multiscale decomposition given by Eq. (14) in Section is:

$$\begin{aligned}
u_{i,j,k}^{n+1} = & f_{i,j,k} \quad (15) \\
& + \frac{1}{2h^2} \left[\frac{u_{i+1,j,k}^n - u_{i,j,k}^{n+1}}{\sqrt{\epsilon^2 + (D_{+x}u_{i,j,k}^n)^2 + (D_{0y}u_{i,j,k}^n)^2 + (D_{0z}u_{i,j,k}^n)^2}} \right] \\
& - \frac{1}{2h^2} \left[\frac{u_{i,j,k}^{n+1} - u_{i-1,j,k}^n}{\sqrt{\epsilon^2 + (D_{-x}u_{i,j,k}^n)^2 + (D_{0y}u_{i-1,j,k}^n)^2 + (D_{0z}u_{i-1,j,k}^n)^2}} \right] \\
& + \frac{1}{2h^2} \left[\frac{u_{i,j+1,k}^n - u_{i,j,k}^{n+1}}{\sqrt{\epsilon^2 + (D_{0x}u_{i,j,k}^n)^2 + (D_{+y}u_{i,j,k}^n)^2 + (D_{0z}u_{i,j,k}^n)^2}} \right] \\
& - \frac{1}{2h^2} \left[\frac{u_{i,j,k}^{n+1} - u_{i,j-1,k}^n}{\sqrt{\epsilon^2 + (D_{0x}u_{i,j-1,k}^n)^2 + (D_{-y}u_{i,j,k}^n)^2 + (D_{0z}u_{i,j-1,k}^n)^2}} \right] \\
& + \frac{1}{2h^2} \left[\frac{u_{i,j,k+1}^n - u_{i,j,k}^{n+1}}{\sqrt{\epsilon^2 + (D_{0x}u_{i,j,k}^n)^2 + (D_{0y}u_{i,j,k}^n)^2 + (D_{+z}u_{i,j,k}^n)^2}} \right] \\
& - \frac{1}{2h^2} \left[\frac{u_{i,j,k}^{n+1} - u_{i,j,k-1}^n}{\sqrt{\epsilon^2 + (D_{0x}u_{i,j,k-1}^n)^2 + (D_{0y}u_{i,j,k-1}^n)^2 + (D_{-z}u_{i,j,k}^n)^2}} \right].
\end{aligned}$$

References

- [1] N. M. ALPERT, J. F. BRADSHAW, D. KENNEDY, AND J. A. CORREIA, *The principal axes transformation – A method for image registration*, Journal of Nuclear Medicine, **31**, no. 10, pp. 1717–1722.
- [2] R. BAJCSY, ET AL., *A computerized system for the elastic matching of deformed radiographic images to idealized atlas images*, Journal of Computer Assisted Tomography, **7**, no. 4, pp. 618–625, 1983.
- [3] F. L. BOOKSTEIN, *Principal Warps: Thin-Plate Splines and the Decomposition of Deformations*, IEEE Transactions on Pattern Analysis and Machine Intelligence, **11**, no. 6, pp. 567–585, 1989.
- [4] L. BROWN, *A Survey of Image Registration Techniques*, ACM Computing Surveys, **24**, no. 4, pp. 325–376, 1992.
- [5] G. E. CHRISTENSEN, R. D. RABBITT, AND M. I. MILLER, *Deformable templates using large deformation kinematics*, IEEE Transactions on Image Processing, **5**, pp. 1435–1447, 1996.
- [6] W. R. CRUM, T. HARTKENS, AND D. L. G. HILL, *Non-rigid image registration: theory and practice*, The British Journal of Radiology, **77**, pp. 140–153, 2004.
- [7] C. DE BOOR, *A Practical Guide to Splines*, Berlin, 1978.
- [8] M. FERRANT, A. NAVAVI, B. MACQ, P. M. BLACK, F. A. JOLESZ, R. KIKINIS, ET AL., *Serial registration of intraoperative MR images of the brain*, Medical Image Analysis, **6**, pp. 337–359, 2002.
- [9] *Insight Segmentation and Registration Toolkit (ITK) Software Guide*, <http://www.itk.org>.
- [10] J. LIAN, L. XING, S. HUNJAN, et al., *Integrating deformable MRI/MSRI and CT image registration into the prostate IMRT treatment planning*, Medical Physics, **31**, no. 11, pp. 3087–3094, 2004.
- [11] W.A. LIGHT, *Variational methods for interpolation, particularly by radial basis functions*, Numerical Analysis. pp. 94–106, 1995.
- [12] J. MODERSITZKI, *Numerical Methods for Image Registration*, Oxford, 2004.
- [13] D. PAQUIN, D. LEVY, E. SCHREIBMANN, AND L. XING, *Multiscale image registration*, Mathematical Biosciences and Engineering, **3**, no. 2, pp. 389–418, 2006.
- [14] D. PAQUIN, D. LEVY, AND L. XING, *Multiscale Deformable Registration*, submitted to IEEE Transactions on Medical Imaging.
- [15] D. W. ROBERTS, A. HARTOV, F.E. KENNEDY, M.I. MIGA, AND K.D. PAULSEN, *Intraoperative brain shift and deformation: A quantitative analysis of cortical displacement in 28 cases*, Neurosurgery, **43**, no. 4, pp. 749–760, 1998.
- [16] K. ROHR, *Landmark-based image analysis*, Computational Imaging and Vision, Kluwer Academic, Dordrecht, 2001.

- [17] L. RUDIN, S. OSHER, AND E. FATEMI, *Nonlinear total variation based noise removal algorithms*, Physica D, **60**, pp. 259–268, 1992.
- [18] E. SCHREIBMANN, G. T.Y. CHEN, AND L. XING, *Image interpolation in 4D CT using a bspline deformable registration model*, International Journal of Radiation Oncology, Biology, Physics, **64**, no. 5, pp. 1537–1550.
- [19] E. SCHREIBMANN AND L. XING, *Narrow band deformable registration of prostate MRI/MRSI and CT studies*, International Journal of Radiation Oncology, Biology, Physics, **62**, no. 2, pp. 595–605, 2005.
- [20] E. SCHREIBMANN AND L. XING, *Image registration with auto-mapped control volumes*, Medical Physics, **33**, no. 4, pp. 1165–1179, 2006.
- [21] E. TADMOR, S. NEZZAR, AND L. VESE, *A multiscale image representation using hierarchical (BV, L^2) decompositions*, Multiscale Modeling and Simulations, **2**, no. 4, pp. 554–579, 2004.
- [22] A. TEI, A. BHARATHA, M. FERRANT, P. McL. BLACK, F.A. JOLESZ, R. KIKINIS, AND S. WARFIELD, *Tracking volumetric brain deformation during image guided neurosurgery*, VISIM: Information Retrieval and Exploration in Large Medical Image Collections, Springer Verlag, 2001.